

# Quantum circuit fidelity estimation using machine learning

Avi Vadali<sup>1</sup>, Rutuja Kshirsagar<sup>2</sup>, Prasanth Shyamsundar<sup>3</sup>, and Gabriel N. Perdue<sup>3</sup>

<sup>1</sup>California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125, USA

<sup>2</sup>Zapata Computing, Inc. 100 Federal Street, FL 20, Boston, MA 02110 USA

<sup>3</sup>Fermi National Accelerator Laboratory, Fermilab Quantum Institute, PO Box 500, Batavia, Illinois, 60510-0500 USA

**The computational power of real-world quantum computers is limited by errors. When using quantum computers to perform algorithms which cannot be efficiently simulated classically, it is important to quantify the accuracy with which the computation has been performed. In this work we introduce a machine-learning-based technique to estimate the fidelity between the state produced by a noisy quantum circuit and the target state corresponding to ideal noise-free computation. Our machine learning model is trained in a supervised manner, using smaller or simpler circuits for which the fidelity can be estimated using other techniques like direct fidelity estimation and quantum state tomography. We demonstrate that the trained model can predict the fidelities of more complicated circuits for which such methods are infeasible.**

## 1 Introduction

Estimating the quality of computations performed by a noisy quantum computer is an important task. It allows us to benchmark current and future quantum computers and establish trust in their computations [15]. Furthermore, estimating the quality of different circuit implementations of a given quantum algorithm can help choose an optimal implementation in a noise-aware manner. The quality of a quantum computation can be quantified using the fidelity between the quantum state produced by the computation and the target quantum state corresponding to an ideal (noise-

less) computation. Henceforth, we will refer to this fidelity (assuming that all the qubits are initialized to  $|0\rangle$  prior to the computation) simply as the fidelity of the circuit. Several techniques exist in the literature to estimate this fidelity, e.g., quantum state tomography [13], direct fidelity estimation [4, 5], Mirror Circuit Fidelity Estimation (MCFE)<sup>1</sup> [15].

The available techniques vary both in a) situations where they are applicable, and b) their classical and quantum computational cost. Techniques like quantum state tomography and direct fidelity estimation can accurately estimate fidelities for all noise models. However, they require access to the target quantum state from classical simulations, and involve running the circuit under consideration multiple times on the quantum hardware. The classical and quantum computational costs required for these techniques can be prohibitively expensive for large circuits.

The MCFE technique does not require access to the target state from simulations. It is compatible with a wide range of noise models and may even account for calibration errors in the implementation of the gates that are exactly cancelled by their inverses through the insertion of a random Pauli layer between the forward and inverse circuits [14]. However, it still requires one to run multiple randomly chosen variations of the relevant circuit on the quantum hardware.

In this paper, we are interested in fidelity estimation techniques which do not require running the relevant circuits on a quantum computer. Such techniques rely upon learning the noise-characteristics of the quantum device in advance through appropriate experiments, and using this to predict the fidelities of new circuits. For convenience, we will refer to such techniques as “fully-

<sup>1</sup>MCFE is a technique for estimating the entanglement or process fidelity, which is related but not identical to state fidelity considered in this paper.

Avi Vadali: [avadali@caltech.edu](mailto:avadali@caltech.edu), This work was completed while the first author was at the Latin School of Chicago.

Rutuja Kshirsagar: [rutuja.kshirsagar@zapatacomputing.com](mailto:rutuja.kshirsagar@zapatacomputing.com), This work was completed while the second author was at Virginia Tech.

classical” fidelity estimation techniques, despite the initial learning phase, which requires quantum hardware. An example of such a technique is estimating the fidelity of a circuit as the product of fidelities of all the gates used in the circuit, which are measured ahead of time during calibration stages. Fully-classical fidelity estimation is extremely useful, e.g., to perform an initial evaluation and screening of the, typically very large number of, potential circuit implementations of a quantum algorithm. This allows a user to narrow the focus down to a smaller number of candidate circuit implementations that may be further refined with more sophisticated, albeit more computationally expensive, techniques like cross-entropy benchmarking (XEB) calibration [1].

However, fully-classical fidelity estimation is challenging to perform accurately. The difficulty of fully-classical fidelity estimation arises from the complex and difficult-to-model nature of noise in quantum devices. For example, it is well known that the product of gate fidelities is a poor predictor of circuit fidelities. We view ML algorithms as a heuristic family of models capable of implicitly fitting a parameterized model of the performance of a quantum computer.

One of the main challenges in deploying ML for the purpose of fidelity estimation is developing a training dataset. Here we propose an approach whereby we utilize small sections of a larger quantum computer to build a library of circuits based purely on Clifford gates. We then show it is possible for a given simulated noise model to use the model for inference on larger circuits with higher degrees of entanglement and for non-Clifford gates. Testing this algorithm for other noise models, in particular for real quantum hardware, is work in progress. We do not claim the ML model deployed here is optimal for this task, only that it is able to perform the task to a useful degree of accuracy.

In this work, we demonstrate the use of machine learning to predict the fidelity of noisy quantum circuits. The idea is to train machine learning (ML) models in a supervised manner, using circuits for which the fidelity estimation can be performed using other techniques like direct fidelity estimation (DFE) and quantum state tomography. Post training, we will use the ML model to predict the fidelity of circuits where such techniques are not feasible. Importantly, once the

model is trained, it allows us to make fidelity estimates without access to the quantum hardware. The crucial step is the construction of appropriate training datasets, which allow for learning the noise characteristics of the quantum device and successful generalization of prediction of fidelity. In particular, we focus on the following two generalizations:

**Clifford to non-Clifford circuits.** The fidelity of Clifford circuits can be estimated efficiently using DFE. The technique however does not work for general (non-Clifford) circuits. Here, we train an ML-based fidelity predictor on Clifford circuits, and use it to predict the fidelity of non-Clifford circuits. In order for the ML model to generalize from Clifford to non-Clifford circuits, the circuits in the training dataset need to contain non-Clifford gates. This is accomplished by introducing non-Clifford gates along with their inverses into Clifford circuits in sequence — we refer to this as the “Clifford-reducible” construction. We hypothesize this technique allows the ML algorithm to learn the noise characteristics of non-Clifford gates while still producing a circuit where the noiseless target state is classically accessible.

**Tiled to non-tiled circuits.** The fidelity of circuits performed on a small number of qubits can be estimated efficiently using quantum state tomography. We leverage this to train our ML models by constructing “tiled” circuits as follows: We construct (factorize) quantum circuits with non-overlapping segments on the qubit lattice (called tiles) with no entangling gates acting across these segments through the whole of the circuit. For such cases, the fidelity of the overall circuit can be computed as the product of the fidelities of the sub-circuits within each tile. By using circuits with different tiling structures in the training dataset, we can expose the ML algorithm to a full sampling of the local noise behaviors and allow for generalization to non-tiled circuits.

There is a feature of the tiling simulation studies in this work that requires mention: it is possible to completely remove inactive qubits. Furthermore, resource constraints and simulation costs form a strong incentive to do so. The noise model presented in this work contained nearest-neighbor interactions, including a cross-talk component modeled as a  $ZZ$  unitary interaction. As a cost-saving measure, we did not include a par-

tial cross-talk effect when simulating tiled circuits as subsamples of the full qubit lattice. Real quantum hardware does not allow for this removal, but in exchange it will naturally model the interactions with qubits outside the tiled sub-circuit.

All of our experiments are performed using simulations with custom models. We utilize the `Cirq`[2] and `Qsimcirq` software packages to simulate noiseless and noisy quantum circuits. Demonstration experiments on hardware devices are planned for future work.

## 2 Previous work

Noise is an irremovable part of quantum circuits on current hardware devices. Multiple studies have been performed relating the domains of machine learning, quantum information, noise in quantum circuits.

Strikis, Qin, Chen, Benjamin and Li [9] studied the process of learning a statistical model to predict output distributions and used machine learning techniques to provide a strategy for quantum error compensation. They provided an efficient protocol for quantum noise mitigation. Proctor, Seritan, Nielsen, Rudinger, Young, Blume-Kohout and Sarovar [15] estimate the fidelity by modifying the algorithm’s quantum circuit to incorporate the error. Flurin, Martin, Hacohen-Gourgy, and Siddiqi [6] train a neural network in real time to give a trajectory of evolution of the superconducting qubit under certain noise models such as decoherence, etc. Harper, Flammia and Wallman [8] give a protocol to estimate quantum noise and determine correlations in subsets of qubits on a 14-qubit superconducting quantum architecture.

Zhang et al [18] utilize machine learning to make predictions about quantum fidelity from a limited set of measurements. While considering a subset of the full set of possible Pauli operators they show it is possible to use a neural network to predict which of a finite set of fidelity intervals the state belongs to. Similar to this work, Zhang et al use machine learning as a heuristic algorithm to compress information about quantum state fidelity, but they are not making predictions about circuit performance. Rather, their work relates to improving the shot efficiency of fidelity estimation measurements given a final state wavefunction, regardless of how it is produced.

In [10] Huang, Chen, and Preskill propose an efficient ML algorithm for predicting the expectation value of an operator over a process, like a quantum circuit of arbitrary depth, within a small error, given query access to the process. Importantly, the algorithm may be trained using only random product input states and randomized Pauli measurements on the output states, but inference is possible with highly-entangled inputs, and for arbitrary expectation values.

Liu and Zhou [12] employ a “black-box” model utilizing physically motivated parameters (e.g., circuit depth, the number of CNOT gates utilized, etc.) and infer a circuit reliability benchmark — the probability of successful trial (PST, [16]) — using a polynomial fit and a shallow neural-network. The PST uses the probability (estimated with Monte Carlo) of the all-zero’s bit-string running a given circuit and its inverse in sequence. Our work differs in that our machine learning model learns an implicit representation of the noise, and we attempt to generalize training based on compartmentalized entanglement to inference tasks on circuits with broader entanglement. We also use a different reliability benchmark, although theirs has the advantage of being directly portable to hardware. We utilize an approach similar to their black-box approach as a benchmark (see the “gate-counting regressor” in Section 4.2).

Wang et al [17] is the closest work to this one, and debuted just before this work. They use a graph transformer NN model to learn the PST using both simulated and real hardware circuits. Their model includes some explicit qubit parameters in the graph model as well as some global parameters, making a hybrid implicit-explicit model. They additionally consider a set of structured circuits as well as random circuits.

In this paper, we use a neural network (NN) to estimate the fidelity of a quantum circuit which the network has not been seen previously. The NN must be trained on circuits run on a specific quantum computer, but once the algorithm has been fit, our approach no longer requires access to quantum resources to make fidelity estimates. The largest difference between this work and prior works is our focus on utilizing quantum circuits that may be simulated on classical resources with low cost for training by using “hidden” non-Clifford gates and compartmentalized

entanglement, and extending this result to inference on circuits that are substantially harder to simulate classically. We consider a different family of neural networks (3D convolutional networks), but consider the main innovation in this paper to be a focus on training with classically simulable circuits with a goal of extending application to beyond-classical circuits. Here we use the exact circuit fidelity for a metric but this is temporary and will not scale to significant hardware circuits. See Section 5 for future plans.

### 3 Methodology

#### 3.1 Data generation

In this work we choose a noise model, adapted from [7], and generate random circuits of fixed depth using the connectivity map of a square lattice quantum computer, with nearest neighbor connectivity for the 2-qubit gates. For tiling purposes we include a variety of rectangular and square grids of varying sizes. For benchmark scoring of our techniques we use  $3 \times 3$  and  $5 \times 5$  qubit square-lattice grids. We did not engineer our dataset to create a particular distribution of fidelities.

**Circuit depth** We consider two basic approaches. First, we study circuits that are only one moment deep. A moment is a set of gates that are all running in the same slice of time (functionally simultaneously) [3]. Models trained on this type of circuit may be used to evaluate multi-moment circuits by multiplying the fidelity of each moment together as an approximation. This approach has an advantage in that it may be easily extended to an arbitrary number of moments, but the disadvantage of assuming the circuit fidelity of each moment is independent of the input wavefunction. Second, we study multi-moment circuits directly. This model is less flexible but requires fewer assumptions about quantum circuit behavior. Note our simulations do not include time dependent noise sources, so the multiplicative factorization is a reasonable assumption. This is not true in general for modern quantum hardware.

**Clifford-reducible circuits** We construct some training circuits with both Clifford and non-Clifford gates in such a way that they are equivalent, in the noiseless case, to Clifford circuits. This is accomplished by introducing ( $T$ ,  $T^\dagger$ ) pairs

in sequence over pairs of moments. These gates combine to create an effective identity operation on the qubit they are applied to. See Figure 1 for an illustration. This allows for the possibility of using direct fidelity estimation to estimate the fidelity of the circuits (by removing the pairs to restore efficiently simulable behavior), which is considerably faster than tomography.

**Distribution of gates** We consider a set of six single-qubit gates:  $X$ ,  $Y$ ,  $Z$ ,  $H$ ,  $T$ ,  $T^\dagger$  (with  $T$  and  $T^\dagger$  appearing only in multi-moment Clifford-reducible circuits); and a set of two two-qubit gates:  $CX$  (CNOT) and  $CZ$ . We choose gates randomly from this set to generate a circuit. The placement of gates in the circuit is discussed later.

**Tiled circuits** Some training circuits for the  $5 \times 5$  grid are generated by generating smaller square and rectangular lattices and combining them together edge-to-edge in a process we refer to as “tiling.” Tiled circuits have separable target states. This allows us to estimate the fidelity of the complete circuit as a product of fidelities of the component circuits. This reduces computation costs for estimating fidelities for the training circuits, in comparison to non-tiled circuits, which may feature higher degrees of entanglement.

**Noise model** We consider two approaches to simulating noise in our quantum circuits. The first, referred to here as the stochastic product approach, utilized a set of models where we labeled circuits as having succeeded or failed stochastically, with failure probabilities assigned to one and two-qubit gate operations in an uncorrelated fashion. In this approach we looped over all the elements in a circuit and computed an overall fidelity proxy by multiplying the gate success probabilities together. We considered a set of varying gate failure probabilities, detailed in Section 4.1. We tested for differing failure rates as a function of location on the chip and as a function of the gate applied to test whether the ML model could discover these different patterns.

In our second approach, we ran full circuit simulations using `Cirq`, including a custom quantum noise model. This simulated noise model utilizes code from [7]. Note that we study both single and multi-moment circuits with full simulation. It implements two kinds of noise: symmetric local depolarization on 1 and 2 qubit gates and crosstalk



applied to CNOT gates. The depolarizing operator acts upon an  $n$ -qubit system as follows:

$$D_n(\rho, \epsilon) = (1 - \epsilon)\rho + \frac{\epsilon}{2^n}I$$

Here,  $\rho$  is the density matrix of the system and  $\epsilon$  is the probability of depolarization occurring. In [7], a formulation of  $D_n(\rho, \epsilon)$  as a sum over Pauli operators is given. Our noise model applies  $D_n(\rho, \epsilon_1)$  after each single qubit gate and  $D_n(\rho, 10\epsilon_1)$  after each 2-qubit gate. For our purposes, we fixed  $\epsilon_1 = 0.1$  so that there is a 0.1% and 1% chance of applying local depolarization after each single-qubit and double-qubit gate respectively. The unitary operator  $U_{ZZ}[\zeta]$  is given by,

$$U_{ZZ}[\zeta] = \exp(-i2\pi\zeta T |11\rangle\langle 11|).$$

The parameter  $T$  dictates the amount of time the crosstalk gate is applied for, and  $\zeta$  determines the strength of the crosstalk noise. To apply this noise model, we insert  $U_{ZZ}[\zeta]$  after each two qubit gate. Our crosstalk strength parameters are  $T = 10^{-8}$  seconds and  $\zeta = 150000$ . When using full circuit simulation, because we have the full state information we change the definition of fidelity to be the true fidelity of the circuit. Note that in our second approach we do not use location-dependent or gate dependent noises (beyond the distinction between one and two-qubit gates) for reasons we will explain shortly.




Figure 1: Sample Clifford-reducible circuit <sup>a</sup>. Each vertical slice of the circuit represents a moment, or “layer” of gates. In this example we have five moments.

<sup>a</sup>Figure generated using the IBM Quantum Composer: <https://quantum-computing.ibm.com>

### 3.2 Circuit representation with one-hot encoding

Our quantum circuits are represented as 4-dimensional tensors. The first dimension encodes

the moment within the circuit:  $C[0]$  is the first layer,  $C[1]$  is the second layer, etc. See Figure 1 for a pictorial representation of a circuit acting on a set of qubits through moments of gates. The second and third dimensions are the Cartesian  $x$  and  $y$  coordinates of a given qubit on the square lattice respectively. Circuits are implemented assuming nearest-neighbor connectivity on a square lattice, so  $C[n][x][y]$  gives the qubit in moment  $n$  at  $(x, y)$ . The fourth dimension encodes a one-hot vector where a 1 at a particular index indicates the specific gate applied to the corresponding qubit. The possible gates are listed in Section 3.1. For instance,  $C[n][x][y][2]$  equals 1 if there is a  $Y$  gate applied to the qubit at  $(x, y)$  in layer  $n$ . For two-qubit gates, the one-hot index also specifies the secondary qubit that the gate acts upon. We convert tensor-formatted circuits to a **Cirq** representation for simulation by iterating through the tensor and applying gates according to the one-hot encoding.

### 3.3 Neural network architecture and training

For single-moment circuit regressor neural networks, we chose to use a 2D Locally Connected Network [11] without any form of regularization or input normalization. For multi-moment circuit regressor neural networks, we used a 3D Convolutional Neural Network (CNN) with no regularization and no input or batch normalization. Additional details on training are provided in Appendix A. The architectures are provided in Appendix B.

## 4 Experiments and results

### 4.1 Single moment circuits

#### 4.1.1 Stochastic product model noise

We employed a variety of uncorrelated stochastic gate fidelity models:

1. In the first approach, single qubit gates are assigned success probabilities of 0.99, and 2-qubit gates are assigned success probabilities of 0.95. This means we assign a 99% (95%) chance of each one (two) qubit gate operating as intended. Otherwise we label the result as an error state.
2. In the next approach, each gate has a unique success probability between 1 and 0.94:  $I =$

1,  $Y = 0.98$ ,  $Z = 0.97$ ,  $H = 0.96$ ,  $CX = 0.95$ ,  $CZ = 0.94$ . Interpretation of these values is as in the previous case.

3. In the next approach, we set the success probability of the outer qubits in the lattice to be higher than that of inner qubits (to simulate a bulk crosstalk-like effect). Specifically, all the single qubit gates have a success probability of 0.99 and all the two qubit gates have a success probability of 0.95 if they are associated with edge qubits. All the single qubit gates have a success probability of 0.97 and all the two qubit gates have a success probability of 0.93 if they are associated with qubits situated off the edge of the lattice.
4. In the final ad-hoc approach, each gate is assigned a unique success probability between 1 and 0.94:  $I = 1$ ,  $Y = 0.98$ ,  $Z = 0.97$ ,  $H = 0.96$ ,  $CX = 0.95$ ,  $CZ = 0.94$ , and to simulate crosstalk, we assign an additional failure “penalty” (0.94) to all qubits adjacent to a gate. This means for a single qubit gate, we check all 4 of its neighboring qubits and multiply the success probability by 0.94 everytime there is a gate associated with the neighboring qubit. More specifically, given a gate on the lattice such that three of its neighboring qubits have a gate, then the success probability is multiplied by  $(0.94)^3$ . Similarly, for two qubit gates we check all six of its nearest neighbors and multiply by 0.94 each time a gate is encountered.

The locally connected neural network model was able to learn all four of these ad-hoc failure models and accurately predict success probability for  $3 \times 3$  single moment circuits, as shown in Figure 2. Given this success, for simplicity we chose to use a uniform noise model (no variations for different qubits in the chip, or different gates) when performing full circuit simulations for multi-moment circuits.

#### 4.1.2 Circuit simulations with custom noise

We next swapped stochastic product models for a noise model incorporating depolarization and crosstalk implemented in a full circuit simulation using `Cirq`. We also shifted from the stochastic product model of fidelity metrics to using




Figure 2: A 2D locally connected model’s performance on single-moment circuits implemented on a  $3 \times 3$  qubit lattice with stochastic product model noises.

$F(\rho, \sigma) = (\text{tr} \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}})^2$  where  $\rho, \sigma$  are density matrices. We trained a locally connected network on 3000 simulated single-moment Clifford circuits on a  $3 \times 3$  qubit lattice, and found a linear correlation between the true and predicted fidelities of a similar test dataset, as shown in Figure 3.




Figure 3: Fidelity estimates for a 2D locally connected NN model trained on 1500  $3 \times 3$  single-moment Clifford circuits and tested on 1000  $3 \times 3$  single-moment Clifford circuits. The structure in the distribution of values is related to the integer number of possible gates, and is dominated by the number of two-qubit gates, which are noisier in our model (and on most real quantum computers).

## 4.2 Multi-moment circuits

Our primary benchmark for multi-moment circuits is a linear regression model that carries location-specific parameters for fidelity penalties

for one and two-qubit gates. We chose this model as a good proxy for the “practitioner’s benchmark”, which is to count one and two-qubit gates and assign an estimated fidelity penalty for each. We extended the model slightly by making it-location aware — in other words, each different qubit and coupler will have an independent penalty in the fit.

We compare this benchmark, referred to as the “gate-counting regressor” model to the neural network described in Appendix B, with architecture parameters specified in Table 2. Moving forward, we shorten our description of the gate-counting regressor model to simply the “regressor” model even though both the neural-network and the linear model are regressors estimating the fidelity.

Our comparison spans three levels of increasing complexity. In all cases we *train* the estimators using Clifford-reducible, tiled circuits. Recall these circuits are constructed such that non-Clifford gates appear only in pairs that sequence to the identity operation, and could therefore be removed at compile time for efficient classical simulation. Furthermore, tiled circuits restrict the area of entanglement across the total lattice to again allow for inexpensive classical simulation of

the final state (factorized) wavefunction. Tiling was also a practical requirement in this work for noisy circuit simulation. Even though our noise model is relatively simple, it is prohibitively expensive to generate large simulated datasets on  $5 \times 5$  lattices on the commercial cloud computing nodes available to us.

#### 4.2.1 Performance evaluation

Qualitative estimator benchmarks at the *inference stage* are summarized in Figure 4. We first consider a case where we benchmark the estimators on Clifford-reducible, tiled circuits. Both estimators show predictive power for the fidelity, although the more sophisticated neural network model shows a tighter spread in the relationship between true and predicted fidelities. We next consider non-Clifford-reducible, but still tiled circuits. Again both estimators show predictive power, but the simpler regressor model shows a systematic bias in its estimations, along with a broader spread in the true vs predicted fidelities. We finally consider non-Clifford-reducible, non-tiled circuits. Both estimators continue to show predictive power, and similar spreads in the true vs predicted fidelity distributions, but the simpler regressor continues to show a bias.




Figure 4: Normalized heatmaps showing the predicted vs true fidelities for a neural network estimator and a gate counting regressor, both trained on tiled, Clifford-reducible circuits. The top, middle, and bottom panels correspond to testing the estimators on a) tiled, Clifford-reducible circuits, b) tiled, non-Clifford-reducible circuits, and c) non-tiled, non-Clifford-reducible circuits, respectively. The plots also show the Kendall-tau coefficient between the predicted and true fidelities in each case.



Loosely inspired by the receiver operating characteristic (ROC) curve, we studied an additional metric that attempted to capture the “usefulness” of the estimators for deciding whether or not the estimator prediction for fidelity may be relied on. We hypothesize one major use case of models like these would be to decide if an experiment on hardware is likely enough to succeed to be worth investing quantum resources. For the question of whether a given circuit has a fidelity at least  $t$ , the true positive rate and true negative rate of a given predictor are given by

$$\text{TPR}(t) = \Pr(F_p \geq t \mid F_t \geq t), \quad (1)$$

$$\text{TNR}(t) = \Pr(F_p < t \mid F_t < t), \quad (2)$$

where  $F_p$  is the predicted fidelity,  $F_t$  is the true fidelity, and  $t$  is a given threshold value for the fidelity. Here we will use half the sum of TPR and TNR as a metric to evaluate the performance of the different fidelity estimators. It produces a quality curve over the range of fidelities of interest somewhat similar to a receiver operating characteristic (ROC) curve in that the area under the curve may be used to establish higher quality performance over a range of operation. A random guessing strategy will give a flat score across fidelity thresholds at 0.5, while a perfect classifier will give a flat score across fidelity thresholds at 1.0. We prefer this metric to mean-squared error (MSE) to summarize performance because MSE can be misleading when integrated over the full test dataset, and for quantum circuits our score directly addresses the question — “is this circuit worth running on this hardware?”

We show the score as a function of fidelity threshold for both estimators in Figure 5. Note that while a better performing model will have higher scores, it is also the case that the shape of the performance curves is impacted by the distribution of fidelities in the underlying dataset. Furthermore, while our score is inspired by the ROC curve, it is not really correct to compare the areas under these curves as an absolute standard given the influence of the underlying distributions. Nonetheless, we feel these curves provide an additional useful comparison between estimators.




Figure 5: Figures showing the metric  $(\text{TPR}(t) + \text{TNR}(t))/2$  for various fidelity thresholds  $t$ , for the neural network estimator and the gate counting estimator, both trained on tiled, Clifford-reducible circuits. The top, middle, and bottom panels correspond to testing the estimators on a) tiled, Clifford-reducible circuits, b) tiled, non-Clifford-reducible circuits, and c) non-tiled, non-Clifford-reducible circuits, respectively. TPR, TNR and their uncertainties were estimated using Wilson’s formulas for binomial-proportion confidence intervals.

## 5 Conclusions and outlook

In this work we proposed a method to perform heuristic estimation of circuit fidelity using machine learning. We see evidence, for specific noise-models studied here, that it is possible to train a ML model to perform this task usefully, and further that it is possible to build a training dataset from easy-to-simulate circuits that generalizes to harder-to-simulate circuits. We do not claim this is possible on real quantum hardware at this time because the features of real quantum noise are different from our simulated noise model in ways that may be significant. For example, noise on real quantum hardware is often non-stationary and non-local. Further, we do not claim the specific ML model evaluated here is optimal for this task, and we note the model utilized in this work is specific to a square-lattice qubit connectivity.

We note that the machine learning model we used here was capable of making predictions on circuits with arbitrary entanglement over the qubits even when trained on circuits with compartmentalized entanglement. We do not yet know if this capability generalizes to other noise models, but it is suggestive of a method for studying how well a quantum computer might simulate circuits in the beyond-classical regime by studying the performance of smaller circuits amenable to classical simulation.

We emphasize this method is one tool among many for studying and predicting the performance of a noisy quantum computer. While this method requires training a model with performance labels that may be expensive to acquire (for example, full process tomography is prohibitively expensive for more than a few qubits), and there are good reasons to believe it will require re-training when faced with non-stationary noise models, once a model is obtained it allows for the estimation of performance metrics with no further use of the quantum computer. This is potentially valuable in compilers or when deciding whether certain experiments are worth testing on a real QPU, where access time may be especially precious.

In future work we will test this approach on real quantum hardware for performance metric prediction. On hardware, full process tomography is prohibitively expensive for even a moderate number of qubits, making other metrics like MCFE an attractive proxy. Additionally, here we consider

circuits of fixed depth, but we are studying extensions to allow for circuits of varying depth.

## 6 Acknowledgements

We thank E. Peters for useful code, comments, and instructions in adapting the quantum noise model employed in this work. We also thank members of the QMLQCF collaboration, in particular A. del Gado, J. Duarte, and J. R. Vlimant for enlightening discussion about this work.

This document was prepared using the resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy (DOE), Office of Science, HEP User Facility. Fermilab is managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE-AC02-07CH11359. R.K., G. P., and P. S. were supported for this work by the DOE/HEP QuantISED program grant “HEP Machine Learning and Optimization Go Quantum,” identification number 0000240323, and P. S. was additionally supported by the DOE/HEP QuantISED program grant “QCCFP-QMLQCF Consortium,” identification number DE-SC0019219.

## References

- [1] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, Jun 2018. ISSN 1745-2481. DOI: [10.1038/s41567-018-0124-x](https://doi.org/10.1038/s41567-018-0124-x). URL <https://doi.org/10.1038/s41567-018-0124-x>.
- [2] Cirq Developers. Cirq, Aug 2021. URL <https://doi.org/10.5281/zenodo.5182845>.
- [3] Cirq Developers. Cirq documentation, 2022. URL <https://quantumai.google/cirq/build/circuits>.
- [4] Marcus P. da Silva, Olivier Landon-Cardinal, and David Poulin. Practical characterization of quantum devices without tomography. *Phys. Rev. Lett.*, 107:210404, Nov 2011. DOI: [10.1103/PhysRevLett.107.210404](https://link.aps.org/doi/10.1103/PhysRevLett.107.210404). URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.210404>.

- [5] Steven T. Flammia and Yi-Kai Liu. Direct fidelity estimation from few pauli measurements. *Phys. Rev. Lett.*, 106:230501, Jun 2011. DOI: [10.1103/PhysRevLett.106.230501](https://doi.org/10.1103/PhysRevLett.106.230501). URL <https://link.aps.org/doi/10.1103/PhysRevLett.106.230501>.
- [6] E. Flurin, L. S. Martin, S. Hacothen-Gourgy, and I. Siddiqi. Using a recurrent neural network to reconstruct quantum dynamics of a superconducting qubit from physical observations. *Phys. Rev. X*, 10:011006, Jan 2020. DOI: [10.1103/PhysRevX.10.011006](https://doi.org/10.1103/PhysRevX.10.011006). URL <https://link.aps.org/doi/10.1103/PhysRevX.10.011006>.
- [7] Erik Gustafson, Burt Holzman, James Kowalkowski, Henry Lamm, Andy C. Y. Li, Gabriel Perdue, Sergio Boixo, Sergei Isakov, Orion Martin, Ross Thomson, Catherine Vollgraff Heidweiller, Jackson Beall, Martin Ganahl, Guifre Vidal, and Evan Peters. Large scale multi-node simulations of  $\mathbb{Z}_2$  gauge theory quantum circuits using google cloud platform, 2021. URL <https://arxiv.org/abs/2110.07482>.
- [8] Robin Harper, Steven T. Flammia, and Joel J. Wallman. Efficient learning of quantum noise. *Nature Physics*, 16(12):1184–1188, aug 2020. DOI: [10.1038/s41567-020-0992-8](https://doi.org/10.1038/s41567-020-0992-8). URL <https://doi.org/10.1038/s41567-020-0992-8>.
- [9] Marcel Hinsche, Marios Ioannou, Alexander Nietner, Jonas Haferkamp, Yihui Quek, Dominik Hangleiter, Jean-Pierre Seifert, Jens Eisert, and Ryan Sweke. Learnability of the output distributions of local quantum circuits, 2021.
- [10] Hsin-Yuan Huang, Sitan Chen, and John Preskill. Learning to predict arbitrary quantum processes, 2022. URL <https://arxiv.org/abs/2210.14894>.
- [11] Keras. [https://keras.io/api/layers/locally\\_connected\\_layers/](https://keras.io/api/layers/locally_connected_layers/). Accessed: 2022-09-26.
- [12] J. Liu and H. Zhou. Reliability modeling of nisq-era quantum computers. In *2020 IEEE International Symposium on Workload Characterization (IISWC)*, pages 94–105, Los Alamitos, CA, USA, oct 2020. IEEE Computer Society. DOI: [10.1109/IISWC50251.2020.00018](https://doi.org/10.1109/IISWC50251.2020.00018). URL <https://doi.ieeecomputersociety.org/10.1109/IISWC50251.2020.00018>.
- [13] J. F. Poyatos, J. I. Cirac, and P. Zoller. Complete characterization of a quantum process: The two-bit quantum gate. *Phys. Rev. Lett.*, 78:390–393, Jan 1997. DOI: [10.1103/PhysRevLett.78.390](https://doi.org/10.1103/PhysRevLett.78.390). URL <https://link.aps.org/doi/10.1103/PhysRevLett.78.390>.
- [14] Timothy Proctor, Kenneth Rudinger, Kevin Young, Erik Nielsen, and Robin Blume-Kohout. Measuring the capabilities of quantum computers. *Nature Physics*, 18(1):75–79, Jan 2022. ISSN 1745-2481. DOI: [10.1038/s41567-021-01409-7](https://doi.org/10.1038/s41567-021-01409-7). URL <https://doi.org/10.1038/s41567-021-01409-7>.
- [15] Timothy Proctor, Stefan Seritan, Erik Nielsen, Kenneth Rudinger, Kevin Young, Robin Blume-Kohout, and Mohan Sarovar. Establishing trust in quantum computations, 2022. URL <https://arxiv.org/abs/2204.07568>.
- [16] Swamit S. Tannu and Moinuddin K. Qureshi. Not all qubits are created equal: A case for variability-aware policies for nisq-era quantum computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '19*, page 987–999, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450362405. DOI: [10.1145/3297858.3304007](https://doi.org/10.1145/3297858.3304007). URL <https://doi.org/10.1145/3297858.3304007>.
- [17] Hanrui Wang, Pengyu Liu, Jinglei Cheng, Zhiding Liang, Jiaqi Gu, Zirui Li, Yongshan Ding, Weiwen Jiang, Yiyu Shi, Xuehai Qian, David Z. Pan, Frederic T. Chong, and Song Han. Quest: Graph transformer for quantum circuit reliability estimation, 2022. URL <https://arxiv.org/abs/2210.16724>.
- [18] Xiaoqian Zhang, Maolin Luo, Zhaodi Wen, Qin Feng, Shengshi Pang, Weiqi Luo, and Xiaoqi Zhou. Direct fidelity estimation of quantum states using machine learning. *Phys. Rev. Lett.*, 127:130503, Sep 2021. DOI: [10.1103/PhysRevLett.127.130503](https://doi.org/10.1103/PhysRevLett.127.130503). URL <https://link.aps.org/doi/10.1103/PhysRevLett.127.130503>.

## A Neural network training details

For our single-moment circuit network regressors, the loss was computed via the mean-squared error (MSE), and the model was optimized using stochastic gradient descent (SGD) with learning-rate = 0.01, momentum = 0.9, and a Nesterov Accelerated Gradient. During training, both the training loss and validation loss converged to  $O(10^{-8})$ . The architecture is provided in Appendix B.1.

For multi-moment circuit network regressors, we used a 3D Convolutional Neural Network (CNN) to predict fidelities. During training, the loss was computed by mean squared error (MSE), and the model was optimized using stochastic gradient descent (SGD) with learning-rate = 0.01, momentum = 0.9, and a Nesterov Accelerated Gradient. For example, the loss and validation loss of the 3D CNN for 12-moment  $3 \times 3$  Clifford circuits converges to  $O(10^{-4})$ , as shown in Figure 6. The architecture is provided in Appendix B.2.




Figure 6: Loss of 3D CNN trained on 15,000 12-moment  $3 \times 3$  Clifford circuits.

## B Neural network architectures

We designed, by-hand, a pair of neural network (NN) architectures for the single and multi-moment problems.

### B.1 Single-moment network structure

The layer-by-layer details for the single-moment NN model — utilizing 2D locally-connected layers — is given in Table 1.

Table 1: The structure summary for our model for single-moment 13-channel  $3 \times 3$  circuits. We use a locally connected network. The zero-padding 2D layer (ZP2D) pads the input with zeros for use in subsequent layers. The first layer re-shapes the input to account for the one-hot gate encoding. The first locally connected 2D layer (LC2D) uses 16 channels with a kernel size of (2,2) and stride length of (1,1). The subsequent LC2D layers use 16, 16, and 1 channel respectively (kernel size of (1,1) each time).

Layer	Output Shape	Params
ZP2D	(5, 5, 13)	0
LC2D	(4, 4, 16)	13568
LC2D	(4, 4, 16)	4352
LC2D	(4, 4, 16)	4352
LC2D	(4, 4, 1)	272
Flatten	(16)	0
Dense	(64)	1088
Dense	(64)	4160
Dense	(1)	65

Total params: 27,857

### B.2 Mutli-moment network structure

The layer-by-layer details for the multi-moment NN model — utilizing 3D convolutional layers — is given in Table 2.

Table 2: The structure summary for our multi-moment model. We use three successive 3D convolutional layers with 50 channels, kernel size of (4,4,4), and stride length of (1,1,1). The input to the convolutional layers are padded with 0s so that the outputs have the same 3D dimensions as the input.

Layer	Output Shape	Params
Conv3D	(12, 5, 5, 50)	48050
Conv3D	(12, 5, 5, 50)	160050
Conv3D	(12, 5, 5, 50)	160050
Flatten	(15000)	0
Dense	(500)	7500500
Dense	(500)	250500
Dense	(500)	250500
Dense	(500)	250500
Dense	(50)	25050
Dense	(1)	51
Dense	(1)	2

Total params: 8,645,253