

Neural network accelerator for quantum control

David Xu* A. Barış Özgüler[◊] Giuseppe Di Guglielmo[◊] Nhan Tran[◊]
 Gabriel N. Perdue[◊] Luca Carloni* Farah Fahim[◊]
 Columbia University* Fermilab[◊]

ABSTRACT

Efficient quantum control is necessary for practical quantum computing implementations with current technologies. Conventional algorithms for determining optimal control parameters are computationally expensive, largely excluding them from use outside of the simulation. Existing hardware solutions structured as lookup tables are imprecise and costly. By designing a machine learning model to approximate the results of traditional tools, a cheaper method for arriving at the control parameters can be produced. Such a model can then be synthesized into a hardware accelerator for use in quantum systems. In the long term, such an accelerator could be used near quantum computational hardware where traditional computers cannot operate, enabling effective quantum control at a reasonable cost in-the-loop at low latencies without incurring large data bandwidths outside of the cryogenic environment.

KEYWORDS

hls4ml, machine learning, neural networks, quantum control, FPGA, low-power, low-latency

1 INTRODUCTION

Quantum computing has great potential but today’s quantum computers are still encumbered by errors that limit their practical usefulness [1, 2]. Quantum control seeks to enable precise manipulation of quantum hardware via control-theoretical approaches [3?–5]. The result would be an increase in performance of the imperfect hardware. It is desirable to apply such techniques in actual quantum-computational environments, necessitating purpose-built control hardware that can function at low temperature [6?–10]. By approximating existing quantum control algorithms using machine learning (ML), we hope to achieve comparable functionality using significantly less computation. The resulting model could then be implemented as a hardware accelerator at reasonable costs [11].

A qubit (quantum bit) is a quantum computing system’s smallest unit of information. Its quantum state can be described as a linear combination of two orthonormal states $|0\rangle$ and $|1\rangle$: $A = \alpha |0\rangle + \beta |1\rangle$. Quantum gates apply specified transformations, changing quantum states accordingly. We desire to create a target gate U which optimally changes quantum state A into $U * A$. Time-varying pulses are used to produce such action. For some gate, $U(\beta)$, where β is a gate parameter, the problem of gate synthesis can be understood as finding optimal pulse parameters α to produce action $U(\beta)$. Devising pulses for a given value of β is computationally expensive because the size of the underlying pulse parameters (the vector α) can be quite large. Furthermore, there are infinitely many possible β (as generally β is selected from some subset of \mathbb{R}). Conventional tools for such calculations include Jupyter, Quandary, and QuTiP [12–14]. These tools utilize optimization techniques such as gradient ascent to discover optimal α .

In this work, a machine learning model will be trained to determine the appropriate α for producing $U(\beta)$ for any β (Fig. 1).




Figure 1: Simplified overview.

The model will be implemented as a hardware accelerator via hls4ml [15]. hls4ml is a toolset that converts pre-trained ML models into accelerators for FPGA or ASIC implementation with low power or ultra-low latency goals. It requires little designer overhead. Furthermore, the integration with the Google QKeras library allows aggressively quantized deep neural networks to be implemented. The final model is highly resource-efficient and sacrifices little to no accuracy. Since the goal of our present work is to produce an accelerator capable of operating within the harsh conditions of a quantum computing environment, it is desirable for such an accelerator to be small to minimize energy dissipation within the low-temperature quantum environments managed by systems of limited cooling capacity. Hardware implementations need to be fast so that quantum computations can be done within the coherence time since noise is a big bottleneck for current quantum devices.

Note that hardware implementations that function as lookup tables mapping from a fixed number of β to the corresponding α do exist [16]. However, such tables are limited by their precision, as only entries for a finite number of β can be stored. Further discussion of the advantages and disadvantages of such systems can be found in Section 4.2.

For this exploration, a simplified scenario is specified as follows: The gate $U(\beta)$ (X-gate) rotates a qubit about the x-axis by angle $-\pi \leq \beta \leq \pi$. We model the qubit as the ground and first excited states of a transmon, a type of superconducting qubit. The pulses are specified by a 20-dimensional vector α . The dimension "20" is due to the B-spline parametrization we follow in Jupyter [12], which is based on the numerical techniques in Refs. [17, 18]. The dimension of alpha is proportional to the multiplication of number of carrier frequencies and number of B-splines (here 1 and 10, respectively). X-gate brings the state (ground state) to a superposition state as shown in Figure 2, which visualizes on a Bloch sphere the qubit being rotated from its initial $|0\rangle$ quantum state.

The accelerator will be evaluated for accuracy, cost, and performance. Accuracy is measured by comparing the ML-model results




Figure 2: Bloch sphere visualization of a rotation transformation starting from the $|0\rangle$ quantum state.

to those of conventional quantum control algorithms. This metric is considered both for the software model and its synthesized hardware implementation, which may perform differently due to quantization. Accuracy evaluation includes both traditional ML measures such as mean squared error (MSE) and domain-specific measures such as quantum gate fidelity. Cost is measured as the FPGA area of the synthesized result. Performance is measured via an analysis of accelerator latency and throughput.

Fidelity is a metric of the overlap between two unitaries (gates), U_1 and U_2 . It can be defined in several ways. We use the definition [19]:

$$F = \frac{M + |\text{tr}(U)|^2}{M(M+1)}, \quad (1)$$

where $U = U_1^\dagger U_2$ and $M = 2$ is the dimensionality of the Hilbert space for single qubit. We use the following definitions for three types of unitaries considered in this paper. The ‘‘Golden’’ gate is the target gate. ‘‘Optimized’’ gate is calculated via pulse optimizer (Juqbox.jl [12] for this paper), which starts from a random pulse and finds the optimal pulse until reaching a target fidelity (calculated as overlap between ‘‘optimized’’ and ‘‘golden’’ gates). The ‘‘Predicted’’ gate is the one calculated from the ML model. We check the overlap between these three types of gates, where U_1 and U_2 are two of these three types (golden, optimized, predicted).

2 MODEL DEVELOPMENT

Our initial dataset consists of 101 samples. Each sample is a rotation angle upon the x-axis of a qubit (β value) and the corresponding quantum state transformations in the form of 20 parameters (α vector) that were generated with Juqbox, i.e., an open-source software package designed to solve quantum optimal control problems in closed systems [12]. Our later analysis has shown that the models perform well for most angles other than $-\pi$. This was because the dataset used the same pulses for the angles $-\pi$ and π , as both transformations produce the same final quantum state (i.e., the transformation for $-\pi$ actually is carried out in the positive direction). Thus we decided to remove from the dataset the samples




Figure 3: The initial model architecture: a multi-layer perceptron with seven hidden layers, 1,040 parameters.

that were associated with $-\pi$. We used a 60-20-20-percent split for training, validation, and test sets.

As a starting point in our neural-architecture exploration, we chose the multi-layer perceptron (MLP) shown in Fig. 3. The model has a total of seven hidden layers and 1,040 trainable parameters. We trained the model in Tensorflow with Adam optimizer and minimized the mean-squared error (MSE) as a loss function. We set a target of 5,000 epochs with an early-stopping callback [20]. The final MSE value that we measured on the test set was 7.908×10^{-8} . Since there is no ‘‘correct’’ value for MSE, but the lower, the better, and a zero value means the model is perfect, we assumed this value as our reference for the rest of the neural-architecture exploration.

We also compared the expected and predicted results of the pulses on the Bloch sphere using the graphical features of the QuTiP framework [14]. Fig. 4 shows the similarities of the results. This visual comparison further proves that the proposed model could learn the relationship between the final angle and the required quantum state transformations.

We adopted the gate fidelity in Eq. 1 as our last tool to quantify the quality of our results further. Figure 5 shows the fidelity when comparing the gates produced by our model, the traditional software optimizer, and mathematically-derived ground truth. Again, it is evident that the ML model performs poorly around $\beta = -\pi$, suggesting that the $-\pi$ row of the training dataset be removed so that the sign of β values is consistent with the direction of the resultant rotation transformation.

Figure 6 shows the same results with the first three β removed. The traditional software optimizer was configured to ensure fidelity greater than 0.999 between its results and the ground truth, as visible in the figure. It is interesting to note how the curve indicating overlap between the ML model outputs and the ground truth appears as a ‘‘smoothed’’ version of the overlap between the ML model outputs and the optimizer. Such a result suggests that the ML model can ignore lots of the random noise from the optimizer while effectively learning how to produce the target gate. The ML model




Figure 4: Comparison of transformations produced by Jqbox optimizer and our machine-learning model. Each row illustrates results for a different angle β .

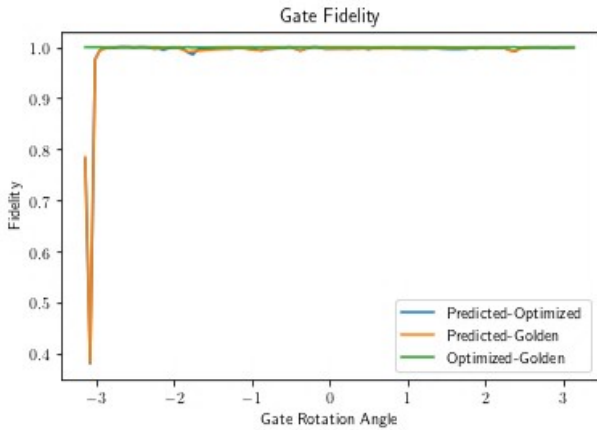


Figure 5: Gate fidelity of forward model. Gates from ML model (“Predicted”), Jqbox optimizer (“Optimized”), and golden mathematically-derived (“Golden”) gates.

produces gates with a gate fidelity of over 0.99 for all remaining β , suggesting practical usefulness.

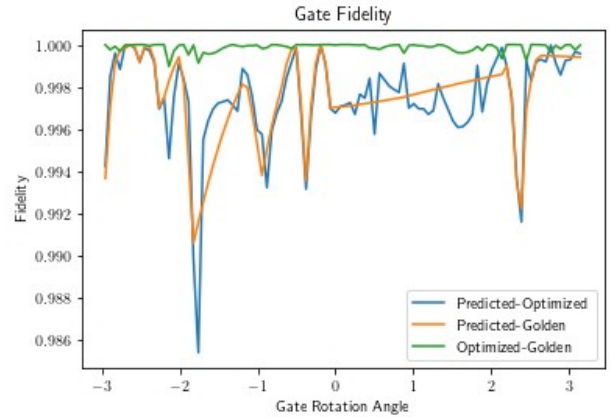


Figure 6: Gate fidelity of forward model. Outlier results for smallest three β removed.

3 MODEL CODESIGN WITH HLS4ML

hls4ml is an open-source framework for the codesign of optimized neural networks and deployment on field-programmable gate arrays (FPGAs) and custom hardware (ASIC) [15]. At its core, hls4ml translates machine-learning models from common open-source software frameworks such as Tensorflow into a register-transfer level (RTL) implementation using high-level synthesis (HLS) tools [21].

The hls4ml framework originates from the Fast Machine Learning for Science community [22], whose focus is the development of tools for scientific applications. The framework includes tools for the design-space exploration and the final FPGA or ASIC implementations. The resulting hardware implementations are configurable, spatial dataflow architectures tailored for speed and efficiency with extreme flexibility in the data-type precision [23].

In hls4ml, a designer can trade off the performance (i.e., latency and throughput) and resource utilization for a model by varying the parallelization of the algorithm via several configuration parameters. For example, the *reuse factor* (RF) parameter controls how many times each multiplier resource is used in the final hardware implementation: a designer with the goal of low latency will choose the lowest RF value.

In this work, we used hls4ml, and additionally, we combined quantization-aware training in QKeras [24] with model- and hardware-centric optimizations to find models that simultaneously accomplish the goals of high accuracy, low latency, and low FPGA-resource usage. We aimed to minimize resources such as onboard FPGA memory (BRAM), digital signal processing (arithmetic) blocks (DSPs), registers (flip-flops, or FFs), and programmable logic (lookup tables, or LUTs).

Our experiments targeted three off-the-shelf FPGA-based development boards with decreasing resources, thus making the code-sign effort increasingly challenging: a Digilent Genesys 2, an Avnet Ultra96-V2, and a Digilent Arty A7-100T.




Figure 7: A smaller multi-layer perceptron with six hidden layers and 783 parameters.

3.1 Initial hardware implementation

The Digilent Genesys 2 board is a development board equipped with an AMD/Xilinx Kintex-7 FPGA chip (xc7k325tffg900-2), which comes with 890 BRAMs, 840 DSPs, 407,600 FFs, and 203,800 LUTs. We passed as input of hls4ml the model developed in Sec. 2 with a target clock period of 5 ns, a reuse factor of 1, and fixed-point precision with overall 16 bits for word width. The resulting implementation was a pipeline with a latency¹ of 35 cycles and initiation interval² of 1. The synthesized model used 54% of the available DSPs and 13% of available LUTs (other resources are not discussed as their utilization remains well below availability in this and all further experiments described herein).

3.2 Shrinking the model

The Avnet Ultra96-V2 is a smaller development board equipped with an AMD/Xilinx Zynq UltraScale MPSoC (xczu3eg-sbva484-1-e). This SoC combines four ARM Cortex A53 cores and programmable logic that offers 432 BRAMs, 360 DSPs, 141,1120 FFs, and 70,560 LUTs. The synthesis of our previous model returned a DSP utilization of 220% and LUT utilization of 38%. Thus, we took additional steps to reduce the resource requirements of our hardware implementation. We removed one hidden layer and some nodes from the remaining hidden layers. The resulting model shown in Fig. 7 had a total of 783 parameters.

To reduce the final resource requirements on FPGA, besides reducing the overall number of model parameters, we could have increased the reuse factor in the hls4ml configuration, but we decided to keep the initial value of one. While a larger reuse factor would significantly decrease the size of the implementation, doing so also decreases parallelism, impacting latency and throughput.

Instead, to address the resource usage, we opted for a more aggressive quantization during the quantization-aware training with QKeras. We analyzed the model inputs and outputs, and for each layer, we chose a fixed-point precision with a 2-bit-integer part and either 10- or 11-bits for the overall word width. The new model trained in QKeras has an MSE of 6.6093×10^{-8} , which is

¹The latency of a pipeline is the time required for one input to pass through the system from start to end.

²The initiation interval of a pipeline is the time necessary for the system to process the next input (or to produce the next output).

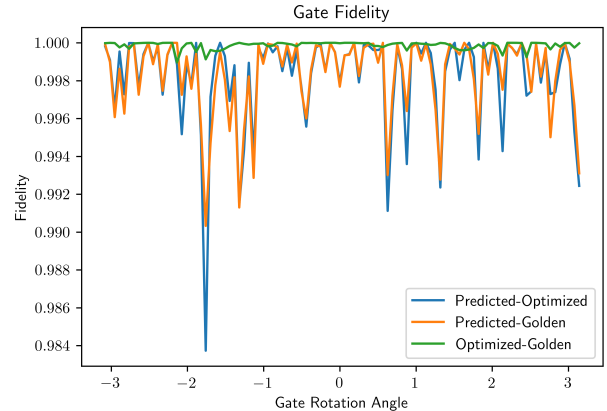


Figure 8: Gate fidelity of the final implementation. Predicted-golden fidelity exceeds 0.99 for all angles (β values).

slightly worse than the original model from Sec.2. But such a change is significant for the underlying FPGA-resource mapping since multipliers on 12 and fewer bits get mapped on LUTs rather than the scarcer DSPs. The resulting model fitted excellently within the Ultra96-V2 constraints: it uses 41% of DSPs and 58% of LUTs.

3.3 Final implementation

The third board we tested is an Arty A7-100T with an AMD/Xilinx Artix-7 chip (xc7a100t-csg324-1). This chip has fewer resources but shows similarities to a future eFPGA platform of interest: 270 BRAMs, 240 DSPs, 126,800 FFs, and 63,400 LUTs.

We adopted the same MLP architecture as in Fig. 7. However, we chose a mixed approach to balance DSP and LUT usage and the overall model accuracy. We used word lengths of 14, 12, 12, and 16 bits (and thus DSPs) for the first three layers and the last layer, respectively, while the word length for the remaining layers was 10 (and thus implemented in LUTs). The number of bits for the integer part was two for all layers, except the output layer, which has zero bits for the integer part. The MSE measured on the test set was 5.7473×10^{-8} .

To fit the board-resource constraints, after the quantization-aware training phase, we manually edited the hls4ml-generated design. We set a smaller 10- and 12-bit-word length for the third and output layer, respectively. Such changes decreased the resource requirements while largely preserving performance. The synthesized implementation fit the Artix-7 chip with a 99% DSP utilization and 62% LUT utilization. The resulting implementation was a pipeline with an initiation interval of 1 and a latency of 35 cycles. Looking at Figure 8, we see that the fidelity between the gates produced by the model and the target gates is above 0.99 for all β .

4 FURTHER CONSIDERATIONS

In this section, we discuss some additional topics and future lines of research.

4.1 Even higher fidelity

Our final model achieves a gate fidelity of over 0.99 for all β examined. However, it is always desirable to increase the fidelity even further. For example, our training data reaches fidelity of over 0.999 for all β , so the question arises: Would it be possible for a model trained on this data to achieve the same level of performance? Various solutions exist that can be combined with our current codesign flow based on hls4ml. We can employ AutoML solutions to automatically discover a more performant model. These approaches are also known as neural architecture search (NAS) [25]. Tools like AutoKeras [26] use a process of searching through neural network architectures to best address the optimal modeling task.

We have also observed that perhaps the function to compute each α differs greatly across the pulse parameters. Thus, we can devise a branched model in which each output is preceded by a hidden layer of its own, with a shared hidden layer at the top connected to the input. We started investigating such a model as in Fig. 9, but we are still in the early phases of the analysis. Similarly, we are investigating larger models (up to 20 times the current number of parameters) that would require larger FPGAs but provide one order for magnitude lower MSE and fidelity that exceeds 0.999, as shown in Fig. 10.

4.2 Implementations with lookup tables

While it is significant that a machine-learning model can achieve comparable accuracy to a specific dataset, a designer may wonder why we cannot simply store the training dataset in a lookup table and use as an index the β value so we can obtain the optimized α values when needed. Admittedly, such an implementation would take up a smaller area and have lower latency for this particular gate.

However, the desired parameter domain may be much larger for more complicated gates. For example, in this initial set of experiments, we have only one gate parameter β , which ranges from $-\pi$ to π . However, consider another gate that performs a rotation to *any* point on the Bloch sphere. The expected number of entries in the lookup table will increase exponentially as a function of the number of input parameters if the same per-parameter granularity is maintained. With an implementation based on neural networks, however, as we are still outputting only 20 pulse parameters, we may expect the network resource requirements to grow slower. For example, one could even imagine that we duplicate the network to produce 20 pulse parameters independently for each input parameter. Then, the designer can use another network to combine these pulse parameters into one final set of pulse parameters. Such a network architecture grows linearly in the number of gate parameters. Furthermore, these networks allow for interpolation and perhaps extrapolation. While a lookup table is constrained to the values stored, the “intelligence” of a neural network allows it to reasonably interpolate the pulse parameters for input gate parameters not seen during training.

Current quantum computing systems implemented using lookup tables will round gate parameters to the closest matching entry in the table [16]. The result is the application of lower-fidelity gates for the desired transformations, as the gate employed will have been optimized for different gate parameters than those of the

desired gate. As a result, the accuracy of algorithmic results will decrease. For example, with the quantum approximate optimization algorithm (QAOA), we would expect the optimization results to be of lower quality, as gates are poorly approximated at each step [27].

5 CONCLUSIONS

To summarize the main contributions of this work, we have developed a neural network capable of predicting pulse parameters for a quantum gate (specifically, a rotation on the Bloch sphere on the x-axis) to the fidelity of 0.99 and shown an FPGA implementation of this network that fits a small off-the-shelf development board. This implementation exhibits a latency of 175 ns and is pipelined with an initiation interval of 1 and a clock period of 5 ns. Our experimental setup is publicly available at github.com/fastmachinelearning/ml4quantum-release.

Preliminary results have shown that defining a machine learning model with a fidelity higher than 0.999 is possible. We leave the FPGA deployment of such a model to future investigation together with the training with further-optimized training data, the augmentation of the dataset with data generated using additional random seeds, the design of a custom eFPGA-based system, the performance evaluation in extreme low-temperature environments, the comparison with lookup-table-based solutions, and the implementation of more complicated gates.

One may imagine that such models could be extended to control multi-qubit systems performing complex tasks. For example, application-specific hardware accelerators could be designed to capture multi-stage domain-specific operations in gates generated on-the-fly. Such abilities would allow for cheaper, more precise, and faster quantum computing systems that will be used to solve the problems of tomorrow. By moving the computation into the colder domains of the quantum computing system, the number of wires leading to the warmer domains can be reduced significantly (as we no longer need wires to transmit α values calculated externally). Such wires are costly in the amount of energy they dissipate within the cold domains which are managed by machines of limited cooling capacity [7].

5.1 Acknowledgements

We thank Matt Otten for discussions on pulse parameters and neural networks. This work was conducted as part of the Spring 2022 research seminar COMS E6868 - Embedded Scalable Platforms at Columbia University. DX thanks the TA Joseph Zuckerman for his help throughout the semester. ABÖ and GNP were partially supported by the DOE/HEP QuantISED program grant “HEP Machine Learning and Optimization Go Quantum,” identification number 0000240323 and the DOE/HEP QuantISED program grant “QCCFP-QMLQCF Consortium,” identification number DE-SC0019219. This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

REFERENCES

- [1] Marco Fellous-Asiani, Jing Hao Chai, Robert S Whitney, Alexia Auffèves, and Hui Khoon Ng. Limitations in quantum computing from resource constraints. *PRX Quantum*, 2(4):040335, 2021.

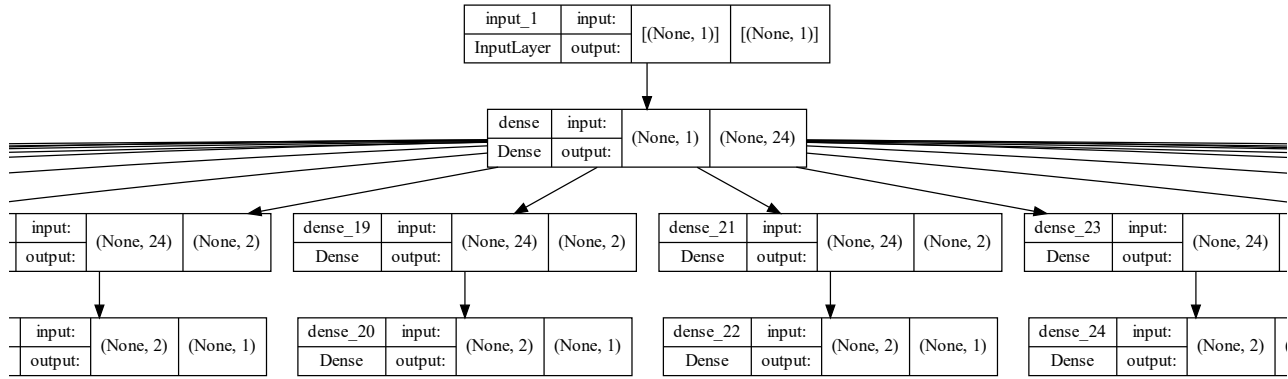


Figure 9: Branched model design. A central crop of the architecture diagram is shown for clarity (the full diagram is available within the documentation repo). The figure continues to the left and right in a similar pattern for all 20 pulse parameters (α).

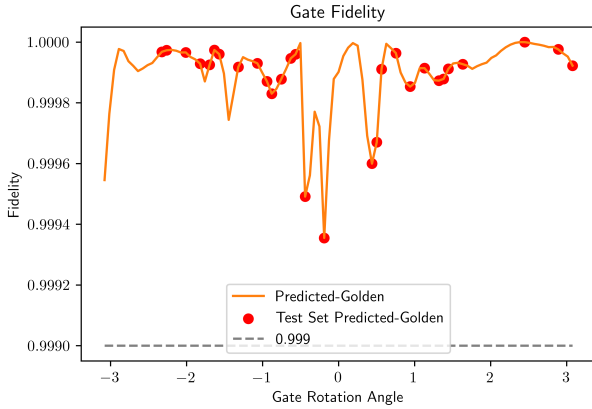


Figure 10: Predicted-golden gate fidelity of large model. Test set points are marked by red circles, showing the ability of the model to interpolate. Fidelity exceeds 0.999 for all test set β .

[2] Harrison Ball, Michael J Biercuk, Andre RR Carvalho, Jiayin Chen, Michael Hush, Leonardo A De Castro, Li Li, Per J Liebermann, Harry J Slatyer, Claire Edmunds, et al. Software tools for quantum control: Improving quantum computer performance through noise and error suppression. *Quantum Science and Technology*, 6(4):044011, 2021.

[3] A. Barış Özgüler and Davide Venturelli. Numerical gate synthesis for quantum heuristics on bosonic quantum processors, 2022.

[4] Christiane P Koch, Ugo Boscain, Tommaso Calarco, Gunther Dirr, Stefan Filipp, Steffen J Glaser, Ronnie Kosloff, Simone Montangero, Thomas Schulte-Herbrüggen, Dominique Sugny, et al. Quantum optimal control in quantum technologies. strategic report on current status, visions and goals for research in europe. *arXiv preprint arXiv:2205.12110*, 2022.

[5] A. Barış Özgüler and Joshua A. Job. Dynamics of qudit gates and effects of spectator modes on optimal control pulses, 2022.

[6] Anastasiia Butko, George Michelogiannakis, Samuel Williams, Costin Iancu, David Donofrio, John Shalf, Jonathan Carter, and Irfan Siddiqi. Understanding quantum control processor capabilities and limitations through circuit characterization. In *2020 International Conference on Rebooting Computing (ICRC)*, pages 66–75. IEEE, 2020.

[7] Xiao Xue, Bishnu Patra, Jeroen PG van Dijk, Nodar Samkharadze, Sushil Subramanian, Andrea Corna, Brian Paquelet Wuetz, Charles Jeon, Farhana Sheikh, Esdras Juarez-Hernandez, et al. Cmos-based cryogenic control of silicon quantum circuits. *Nature*, 593(7858):205–210, 2021.

[8] G Koolstra, N Stevenson, S Barzili, L Burns, K Siva, S Greenfield, W Livingston, A Hashim, RK Naik, JM Kreikebaum, et al. Monitoring fast superconducting qubit dynamics using a neural network. *Physical Review X*, 12(3):031017, 2022.

[9] Leandro Stefanazzi, Kenneth Treptow, Neal Wilcer, Chris Stoughton, Collin Bradford, Sho Uemura, Silvia Zorzetti, Salvatore Montella, Gustavo Cancelo, Sara Sussman, et al. The qick (quantum instrumentation control kit): Readout and control for qubits and detectors. *Review of Scientific Instruments*, 93(4):044709, 2022.

[10] M Sohaib Alam, Sergey Belomestnykh, Nicholas Bommman, Gustavo Cancelo, Yu-Chiu Chao, Mattia Checchin, Vinh San Dinh, Anna Grassellino, Erik J Gustafson, Roni Harnik, et al. Quantum computing hardware for hep algorithms and sensing. *arXiv preprint arXiv:2204.08605*, 2022.

[11] Ramon WJ Overwater, Masoud Babaie, and Fabio Sebastiano. Neural-network decoders for quantum error correction using surface codes: A space exploration of the hardware cost-performance tradeoffs. *IEEE Transactions on Quantum Engineering*, 3:1–19, 2022.

[12] LLNL. Juqbox.jl. URL <https://software.llnl.gov/juqbox.jl/>.

[13] Stefanie Günther, N. Anders Petersson, and Jonathan L. Dubois. Quandary: An open-source C++ package for high-performance optimal control of open quantum systems, 2021.

[14] J.R. Johansson, P.D. Nation, and Franco Nori. Qutip: An open-source python framework for the dynamics of open quantum systems. *Computer Physics Communications*, 183(8):1760–1772, Aug 2012. ISSN 0010-4655.

[15] Farah Fahim, Benjamin Hawks, Christian Herwig, James Hirschauer, Sergio Jindariani, Nhan Tran, Luca P. Carloni, Giuseppe Di Guglielmo, Philip C. Harris, Jeffrey D. Krupa, Dylan S. Rankin, Manuel Blanco Valentin, Josiah D. Hester, Yingyi Luo, John Mamish, Seda Orgrenci-Memik, Thea Aarrestad, Hamza Javed, Vladimir Loncar, Maurizio Pierini, Adrian Alan Pol, Sioni Summers, Javier M. Duarte, Scott Hauck, Shih-Chieh Hsu, Jennifer Ngadiuba, Mia Liu, Duc Hoang, Edward Kreinar, and Zhenbin Wu. hls4ml: An open-source codesign workflow to empower scientific low-power machine learning devices. *CoRR*, abs/2103.05579, 2021. URL <https://arxiv.org/abs/2103.05579>.

[16] Zhiding Liang, Hanrui Wang, Jinglei Cheng, Yongshan Ding, Hang Ren, Zhengqi Gao, Xuehai Qian, Song Han, Weiwen Jiang, and Yiyu Shi. Variational quantum pulse learning, 2022. URL <https://arxiv.org/abs/2203.17267>.

[17] N Anders Petersson, Fortino M Garcia, Austin E Copeland, Ylva L Rydin, and Jonathan L DuBois. Discrete adjoints for accurate numerical optimization with application to quantum control. *arXiv preprint arXiv:2001.01013*, 2020.

[18] N Anders Petersson and Fortino Garcia. Optimal control of closed quantum systems via b-splines with carrier waves. *arXiv preprint arXiv:2106.14310*, 2021.

[19] Line Hjortshøj Pedersen, Niels Martin Møller, and Klaus Mølmer. Fidelity of quantum operations. *Physics Letters A*, 367(1):47–51, 2007. ISSN 0375-9601. URL <https://www.sciencedirect.com/science/article/pii/S0375960107003271>.

[20] Lutz Prechelt. Early stopping-but when? In *Neural Networks: Tricks of the trade*, pages 55–69. Springer, 1998.

[21] Razvan Nane, Vlad-Mihai Sima, Christian Pilato, Jongsok Choi, Blair Fort, Andrew Canis, Yu Ting Chen, Hsuan Hsiao, Stephen Brown, Fabrizio Ferrandi, et al.

- A survey and evaluation of fpga high-level synthesis tools. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 35(10):1591–1604, 2015.
- [22] Allison McCarn Deiana, Nhan Tran, Joshua Agar, Michaela Blott, Giuseppe Di Guglielmo, Javier Duarte, Philip Harris, Scott Hauck, Mia Liu, Mark S. Neubauer, Jennifer Ngadiuba, Seda Ogrenci-Memik, Maurizio Pierini, Thea Aarrestad, Steffen Bähr, Jürgen Becker, Anne-Sophie Berthold, Richard J. Bonventre, Tomás E. Müller Bravo, Markus Diefenthaler, Zhen Dong, Nick Fritzsche, Amir Gholami, Ekaterina Govorkova, Dongning Guo, Kyle J. Hazelwood, Christian Herwig, Babar Khan, Sehoon Kim, Thomas Klijnsma, Yaling Liu, Kin Ho Lo, Tri Nguyen, Gianantonio Pezzullo, Seyedramin Rasoulizhad, Ryan A. Rivera, Kate Scholberg, Justin Selig, Sougata Sen, Dmitri Strukov, William Tang, Savannah Thais, Kai Lukas Unger, Ricardo Vilalta, Belina von Krosigk, Shen Wang, and Thomas K. Warburton. Applications and Techniques for Fast Machine Learning in Science. *Frontiers in Big Data*, 5, 2022.
- [23] Hendrik Borrás, Giuseppe Di Guglielmo, Javier Duarte, Nicolò Ghielmetti, Ben Hawks, Scott Hauck, Shih-Chieh Hsu, Ryan Kastner, Jason Liang, Andres Meza, Jules Muhizi, Tai Nguyen, Rushil Roy, Nhan Tran, Yaman Umuroglu, Olivia Weng, Aidan Yokuda, and Michaela Blott. Open-source FPGA-ML codesign for the MLPerf Tiny Benchmark, 2022. URL <https://arxiv.org/abs/2206.11791>.
- [24] Claudionor N. Coelho, Aki Kuusela, Shan Li, Hao Zhuang, Thea Aarrestad, Vladimir Loncar, Jennifer Ngadiuba, Maurizio Pierini, Adrian Alan Pol, and Sioni Summers. Automatic heterogeneous quantization of deep neural networks for low-latency inference on the edge for particle detectors. 2020. URL <https://arxiv.org/abs/2006.10159>.
- [25] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *The Journal of Machine Learning Research*, 20(1):1997–2017, 2019.
- [26] Haifeng Jin, Qingquan Song, and Xia Hu. Auto-keras: An efficient neural architecture search system. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1946–1956. ACM, 2019.
- [27] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.