

# EXPONENTIAL QUANTUM ADVANTAGE IN COMMUNICATION FOR DISTRIBUTED LEARNING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Training and inference with large machine learning models that far exceed the memory capacity of individual devices necessitates the design of distributed architectures, forcing one to contend with communication constraints. We present a framework for distributed computation over a quantum network in which data is encoded into specialized quantum states. We prove that for certain models within this framework, inference and training using gradient descent can be performed with exponentially less communication compared to their classical analogs, and with relatively modest time and space complexity overheads relative to standard gradient-based methods. To our knowledge, this is the first example of exponential quantum advantage for a generic class of machine learning problems with dense classical data that holds regardless of the data encoding cost. Moreover, we show that models in this class can encode highly nonlinear features of their inputs, and their expressivity increases exponentially with model depth. We also find that, interestingly, the communication advantage nearly vanishes for simpler linear classifiers. These results can be combined with natural privacy advantages in the communicated quantum states that limit the amount of information that can be extracted from them about the data and model parameters. Taken as a whole, these findings form a promising foundation for distributed machine learning over quantum networks.

## 1 INTRODUCTION

As the scale of the datasets and parameterized models used to perform computation over data continues to grow Kaplan et al. (2020); Hoffmann et al. (2022), distributing workloads across multiple devices becomes essential for enabling progress. The choice of architecture for large-scale training and inference must not only make the best use of computational and memory resources, but also contend with the fact that communication may become a bottleneck Pope et al. (2022). When using modern optical interconnects, classical computers communicate by sending bits represented by optical light. This however does not fully utilize the potential of the physical substrate as a communication resource. Given suitable computational capabilities and algorithms, the quantum nature of light can be harnessed as a powerful resource. Here we show that for a broad class of parameterized models, if quantum bits (*qubits*) are communicated instead of classical bits, an exponential reduction in the communication required to perform inference and gradient-based training can be achieved. This protocol additionally guarantees improved privacy of both the user data and model parameters through natural features of quantum mechanics, without the need for additional cryptographic or privacy protocols. To our knowledge, this is the first example of generic, exponential quantum advantage on problems that occur naturally in the training and deployment of large machine learning models. These types of communication advantages help scope the future roles and interplay between quantum and classical communication for distributed machine learning.

Quantum computers promise dramatic speedups across a number of computational tasks, with perhaps the most prominent example being the ability revolutionize our understanding of nature by enabling the simulation of quantum systems, owing to the natural similarity between quantum computers and the world Feynman (1982); Lloyd (1996). However, much of the data that one would like to compute with in practice seems to come from an emergent classical world rather than directly exhibiting quantum properties. While there are some well-known examples of exponential quantum speedups focusing on classical problems, most famously factoring Shor (1994) and related hidden subgroup problems Childs & van Dam (2008), these tend to be isolated and at times difficult to relate to practical applications. For example, even though significant speedups are known for certain ubiquitous problems in machine learning such as matrix inversion Harrow et al. (2009) and principal component analysis Lloyd et al. (2014), the advantage is often lost when including the cost of loading classical data into the quantum computer or of reading out the result into classical

memory Aaronson (2015). In applications where an efficient data access model avoids the above pitfalls, the complexity of quantum algorithms tends to depend on condition numbers of matrices which scale with system size in a way that reduces or even eliminates any quantum advantage Montanaro & Pallister (2015). It is worth noting that much of the discussion about the impact of quantum technology on machine learning has focused on computational advantage. However quantum resources are not only useful in reducing computational complexity — they can also provide an advantage in communication complexity, enabling exponential reductions in communication for some problems Raz (1999); Bar-Yossef et al. (2008). Inspired by these results, we study a setting where quantum advantage in communication is possible across a wide class of machine learning models. This advantage holds without requiring any sparsity assumptions or elaborate data access models such as QRAM Giovannetti et al. (2008).

We focus on compositional distributed learning, known as *pipelining* Huang et al. (2018); Barham et al. (2022). While there are a number of strategies for distributing machine learning workloads that are influenced by the requirements of different applications and hardware constraints Xu et al. (2021); Jouppi et al. (2023), splitting up a computational graph in a compositional fashion (Figure 1) is a common approach. We describe distributed, parameterized quantum circuits that can be used to perform inference over data when distributed in this way, and can be trained using gradient methods (Section 2.1). The ideas we present can also be used to optimize models that use data parallelism as well. In principle, such circuits could be implemented on quantum computers that are able to communicate quantum states.

For the general class of distributed, parameterized quantum circuits that we study, we show the following:

- Even for simple circuits in this class, there is an exponential quantum advantage in communication for the problem of estimating the loss and the gradients of the loss with respect to the parameters (Section 3). This additionally implies a privacy advantage from Holevo’s bound (Section 6).
- For a subclass of these circuits, there is an exponential advantage in communication for the entire training process, not just gradient estimation. This subclass includes circuits for fine-tuning using pre-trained features, and the proof is constructed by through convergence rates under convexity assumptions (Section 4).
- The ability to interleave multiple unitaries encoding nonlinear features of data enables expressivity to grow exponentially with depth, and universal function approximation in some settings implying these models are highly expressive in contrast to popular belief about linear restrictions in quantum neural networks (Section 5).

## 2 PRELIMINARIES

### 2.1 LARGE-SCALE LEARNING PROBLEMS AND DISTRIBUTED COMPUTATION

Pipelining is a commonly used method of distributing a machine learning workload, in which different layers of a deep model are allocated distinct hardware resources Huang et al. (2018); Narayanan et al. (2019). Training and inference then require communication of features between nodes. Pipelining enables flexible changes to the model architecture in a task-dependent manner, since subsets of a large model can be combined in an adaptive fashion to solve many downstream tasks. Additionally, pipelining allows sparse activation of a subset of a model required to solve a task, and facilitates better use of heterogeneous compute resources since it does not require storing identical copies of a large model. The potential for large models to be easily fine-tuned to solve multiple tasks is well-known Brown et al. (2020); Bommasani et al. (2021), and pipelined architectures which facilitate this are the norm in the latest generation of large language models Rasley et al. (2020); Barham et al. (2022). Data parallelism, in contrast, involves storing multiple copies of the model on different nodes, training each on a subsets of the data and exchanging information to synchronize parameter updates. In practice, different parallelization strategies are combined in order to exploit trade-offs between latency and throughput in a task-dependent fashion Xu et al. (2021); Jouppi et al. (2023); Pope et al. (2022).

### 2.2 COMMUNICATION COMPLEXITY

Communication complexity Yao (1979); Kushilevitz & Nisan (2011); Rao & Yehudayoff (2020) is the study of distributed computational problems using a cost model that focuses on the communication required between players rather than the time or computational complexity. It is naturally related to the study of the space complexity of streaming algorithms Roughgarden (2015). The key object of study in this area is the tree induced by a communication protocol whose nodes enumerate all possible communication histories and whose leaves correspond to the outputs of

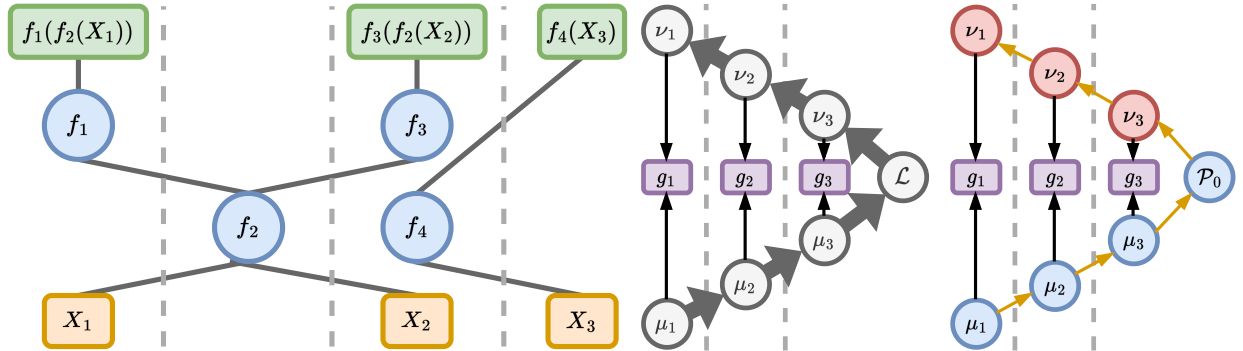


Figure 1: *Left*: Distributed, compositional computation. Dashed lines separate devices with computational and storage resources. The circular nodes represent parameterized functions that are allocated distinct hardware resources and are spatially separated, while the square nodes represent data (yellow) and outputs corresponding to different tasks (green). The vertical axis represents time. This framework of hardware allocation enables flexible modification of the model structure in a task-dependent fashion. *Right*: Computation of gradient estimators  $g_\ell$  at different layers of a model distributed across multiple devices by pipelining. Computing forward features  $\mu_\ell$  and backwards features  $\nu_\ell$  (also known as computing a forward or backward pass) requires a large amount of classical communication (grey) but an exponentially smaller amount of quantum communication (yellow).  $\mathcal{L}$  is the classical loss function, and  $\mathcal{P}_0$  an operator whose expectation value with respect to a quantum model gives the analogous loss function in the quantum case.

the protocol. The product structure induced on the leaves of this tree as a function of the inputs allows one to bound the depth of the tree from below, which gives an unconditional lower bound on the communication complexity. The power of replacing classical bits of communication with qubits has been the subject of extensive study Chi-Chih Yao (1993); Brassard (2001); Buhrman et al. (2009). For certain problems such as Hidden Matching Bar-Yossef et al. (2008) and a variant of classification with deep linear models Raz (1999) an exponential quantum communication advantage holds, while for other canonical problems such as Disjointness only a polynomial advantage is possible Razborov (2002).

At a glance, the development of networked quantum computers may seem much more challenging than the already herculean task of building a fault tolerant quantum computer. However, for some quantum network architectures, the existence of a long-lasting fault tolerant quantum memory as a quantum repeater, may be the enabling component that lifts low rate shared entanglement to a fully functional quantum network Munro et al. (2015), and hence the timelines for small fault tolerant quantum computers and quantum networks may be more coincident than widely believed. As such, it is well motivated to consider potential communication advantages alongside computational advantages when talking about the applications of fault tolerant quantum computers. In Appendix G we briefly survey approaches to implementing quantum communication in practice, and the associated challenges.

In addition, while we largely restrict ourselves here to discussions of communication advantages, and most other studies focus on purely computational advantages, there may be interesting advantages at their intersection. For example, it is known that no quantum state built from a simple (or polynomial complexity) circuit can confer an exponential communication advantage, however states made from simple circuits can be made computationally difficult to distinguish Ji et al. (2018). Hence the use of quantum pre-computation Huggins & McClean (2023) and communication may confer advantages even when traditional computational and communication cost models do not admit such advantages do to their restriction in scope.

### 3 DISTRIBUTED LEARNING WITH QUANTUM RESOURCES

In this work we focus on parameterized models that are representative of the most common models used and studied today in quantum machine learning, sometimes referred to as quantum neural networks McClean et al. (2015); Farhi & Neven (2018); Cerezo et al. (2020); Schuld et al. (2020). We will use the standard Dirac notation of quantum mechanics throughout. A summary of relevant notation and the fundamentals of quantum mechanics is provided in Appendix A. We define a class models with parameters  $\Theta$ , taking an input  $x$  which is a tensor of size  $N$ . The models take the following general form:

**Definition 3.1.**  $\{A_\ell(\theta_\ell^A, x)\}, \{B_\ell(\theta_\ell^B, x)\}$  for  $\ell \in \{1, \dots, L\}$  are each a set of unitary matrices of size  $N' \times N'$  for some  $N'$  such that  $\log N' = O(\log N)$ <sup>1</sup>. The  $\theta_\ell^A, \theta_\ell^B$  are vectors of  $P$  parameters each. For every  $\ell, i$ , we assume that  $\frac{\partial A_\ell}{\partial \theta_{\ell i}^A}$  is anti-hermitian up to a real scaling factor and has at most two eigenvalues, and similarly for  $B_\ell$ .

The model we consider is defined by

$$|\varphi(\Theta, x)\rangle \equiv \left( \prod_{\ell=L}^1 A_\ell(\theta_\ell^A, x) B_\ell(\theta_\ell^B, x) \right) |\psi(x)\rangle, \quad (3.1)$$

where  $\psi(x)$  is a fixed state of  $\log N'$  qubits.

The loss function is given by

$$\mathcal{L}(\Theta, x) \equiv \langle \varphi(\Theta, x) | \mathcal{P}_0 | \varphi(\Theta, x) \rangle, \quad (3.2)$$

where  $\mathcal{P}_0$  is a Pauli matrix that acts on the first qubit.

In standard linear algebra notation, the output of the model is a unit norm  $N'$ -dimensional complex vector  $\varphi_L$ , defined recursively by

$$\varphi_0 = \psi(x), \quad \varphi_\ell = A_\ell(\theta_\ell^A, x) B_\ell(\theta_\ell^B, x) \varphi_{\ell-1}, \quad (3.3)$$

where the entries of  $\varphi_L$  are represented by the amplitudes of a quantum state. The loss takes the form  $\mathcal{L}(\Theta, x) = (\varphi_L^*)^T \mathcal{P}_0 \varphi_L$ , which includes the standard  $L^2$  loss as a special case.

Subsequently we omit the dependence on  $x$  and  $\Theta$  (or subsets of it) to lighten notation, and consider special cases where only subsets of the unitaries depend on  $x$ , or where the unitaries take a particular form and may not be parameterized. Denote by  $\nabla_{A(B)} \mathcal{L}$  the entries of the gradient vector that correspond to the parameters of  $\{A_\ell\}(\{B_\ell\})$ .

While seemingly stringent, the condition on the derivatives is in fact satisfied by many of the most common quantum neural network architectures Cerezo et al. (2020); Crooks (2019); Schuld et al. (2020). This condition is satisfied for example if

$$A_\ell = \prod_{j=1}^P e^{i\alpha_{\ell j}^A \theta_{\ell j}^A \mathcal{P}_{\ell j}^A} \quad (3.4)$$

and the  $\mathcal{P}_{\ell j}^A$  are both unitary and Hermitian (e.g. Pauli matrices), while  $\beta_{\ell j}^A$  are scalars. Such models, or parameterized quantum circuits, are naturally amenable to implementation on quantum devices, and for  $P = O(N^2)$  any unitary over  $\log N'$  qubits can be written in this form. In the special case where  $x$  is a unit norm  $N$ -dimensional vector, a simple choice of  $|\psi(x)\rangle$  is the amplitude encoding of  $x$ , given by

$$|\psi(x)\rangle = |x\rangle = \sum_{i=0}^{N-1} x_i |i\rangle. \quad (3.5)$$

However, despite its exponential compactness in representing the data, a naive implementation of the simplest choice is restricted to representing quadratic features of the data that can offer no substantial quantum advantage in a learning task Huang et al. (2021), so the choice of data encoding is critical to the power of a model. The interesting parameter regime for classical data and models is one where  $N, P$  are large, while  $L$  is relatively modest. For general unitaries  $P = O(N^2)$ , which matches the scaling of the number of parameters in fully-connected networks. When the input tensor  $x$  is a batch of datapoints,  $N$  is equivalent to the product of batch size and input dimension.

The model in Definition 3.1 can be used to define distributed inference and learning problems by dividing the input  $x$  and the parameterized unitaries between two players, Alice and Bob. We define their respective inputs as follows:

$$\begin{aligned} \text{Alice : } & |\psi(x)\rangle, \{A_\ell\}, \\ \text{Bob : } & \{B_\ell\}. \end{aligned} \quad (3.6)$$

The problems of interest require that Alice and Bob compute certain joint functions of their inputs. As a trivial base case, it is clear that in a communication cost model, all problems can be solved with communication cost at most the size of the inputs times the number of parties, by a protocol in which each party sends its inputs to all others. We will be interested in cases where one can do much better by taking advantage of quantum communication.

<sup>1</sup>We will consider some cases where  $N' = N$ , but will find it helpful at times to encode nonlinear features of  $x$  in these unitaries, in which case we may have  $N' > N$ .




Figure 2: Distributed quantum circuit implementing  $\mathcal{L}$  for  $L = 2$ . Both  $\mathcal{L}$  and its gradients with respect to the parameters of the unitaries can be estimated with total communication that is logarithmic in the number of amplitudes  $N$  and the number of parameters per unitary  $P$ .

Given the inputs eq. (3.6), we will be interested chiefly in the two problems specified below.

**Problem 1** (Distributed Inference). *Alice and Bob each compute an estimate of  $\langle \varphi | \mathcal{P}_0 | \varphi \rangle$  up to additive error  $\varepsilon$ .*

The straightforward algorithm for this problem, illustrated in fig. 2, requires  $L$  rounds of communication. The other problem we consider is the following:

**Problem 2** (Distributed Gradient Estimation). *Alice computes an estimate of  $\nabla_A \langle \varphi | \mathcal{P}_0 | \varphi \rangle$ , while Bob computes an estimate of  $\nabla_B \langle \varphi | Z_0 | \varphi \rangle$ , up to additive error  $\varepsilon$  in  $L^\infty$ .*

### 3.1 COMMUNICATION COMPLEXITY OF INFERENCE AND GRADIENT ESTIMATION

We show that inference and gradient estimation are achievable with a logarithmic amount of quantum communication, which will represent an exponential improvement over the classical cost for some cases:

**Lemma 1.** *Problem 1 can be solved by communicating  $O(\log N)$  qubits over  $O(L/\varepsilon^2)$  rounds.*

Proof: Appendix B.

**Lemma 2.** *Problem 2 can be solved with probability greater than  $1 - \delta$  by communicating  $\tilde{O}(\log N (\log P)^4 \log(L/\delta)/\varepsilon^4)$  qubits over  $O(L^2)$  rounds. The time and space complexity of the algorithm is  $\sqrt{P} L \text{poly}(N, \log P, \varepsilon^{-1}, \log(1/\delta))$ .*

Proof: Appendix B.

This upper bound is obtained by simply noting that the problem of gradient estimation at every layer can be reduced to a shadow tomography problem:

**Theorem 1** (Shadow Tomography Aaronson (2017b)). *For an unknown state  $|\psi\rangle$  of  $\log N$  qubits, given  $K$  known two-outcome measurements  $E_i$ , there is an explicit algorithm that takes  $|\psi\rangle^{\otimes k}$  as input, where  $k = \tilde{O}(\log^4 K \log N \log(1/\delta)/\varepsilon^4)$ , and produces estimates of  $\langle \psi | E_i | \psi \rangle$  for all  $i$  up to additive error  $\varepsilon$  with probability greater than  $1 - \delta$ .  $\tilde{O}$  hides a factor  $\text{poly}(\log \log K, \log \log N, \log 1/\varepsilon)$ .*

Using immediate reductions from known problems in communication complexity, we can show that the amount of classical communication required to solve these problem is polynomial in the size of the input, and additionally give a lower bound on the number of rounds of communication required by any quantum or classical algorithm:

**Lemma 3.** *i) The classical communication complexity of Problem 1 and Problem 2 is  $\Omega(\max(\sqrt{N}, L))$ .*

*ii) Any algorithm (quantum or classical) for Problem 1 or Problem 2 requires either  $\Omega(L)$  rounds of communication or  $\Omega(N/L^4)$  qubits (or bits) of communication.*

Proof: Appendix B

The implication of the second result in Lemma 3 is that  $\Omega(L)$  rounds of communication are necessary in order to obtain an exponential communication advantage for small  $L$ , since otherwise the number of qubits of communication required can scale linearly with  $N$ .

Combining Lemma 2 and Lemma 3, in the regime where  $L = O(\text{polylog}(N))$ , which is relevant for classical machine learning models, we obtain an exponential advantage in communication complexity for both inference and gradient estimation. The required overhead in terms of time and space is only polynomial when compared to the straightforward classical algorithms for these problems.

The distribution of the model as in eq. (3.6) is an example of pipelining. Data parallelism is another common approach to distributed machine learning in which subsets of the data are distributed to identical copies of the model. In Appendix C we show that it can also be implemented using quantum circuits, which can then trained using gradient descent requiring quantum communication that is logarithmic in the number of parameters and input size.

Quantum advantage is possible in these problems because there is a bound on the complexity of the final output, whether it be correlated elements of the gradient up to some finite error or the low-dimensional output of a model. This might lead one to believe that whenever the output takes such a form, encoding the data in the amplitudes of a quantum state will trivially give an exponential advantage in communication complexity. We show however that the situation is slightly more nuanced, by considering the problem of inference with a linear model:

**Lemma 4.** *For the problem of distributed linear classification, there can be no exponential advantage in using quantum communication in place of classical communication.*

The precise statement and proof of this result are presented in Appendix E. This result also highlights that the worst case lower bounds such as Lemma 3 may not hold for circuits with certain low-dimensional or other simplifying structure.

## 4 EXPONENTIAL ADVANTAGES IN END-TO-END TRAINING

So far we have discussed the problems of inference and estimating a single gradient vector. It is natural to also consider when these or other gradient estimators can be used to efficiently solve an optimization problem (i.e. when the entire training processes is considered rather than a single iteration). Applying the gradient estimation algorithm detailed in Lemma 2 iteratively gives a distributed stochastic gradient descent algorithm which we detail in Algorithm 2, yet one may be concerned that a choice of  $\varepsilon = O(\log N)$  which is needed to obtain an advantage in communication complexity will preclude efficient convergence. Here we present a simpler algorithm that requires a single quantum measurement per iteration, and can provably solve certain convex problems efficiently, as well as an application of shadow tomography to fine-tuning where convergence can be guaranteed, again with only logarithmic communication cost. In both cases, there is an exponential advantage in communication even when considering the entire training process.

### 4.1 “SMOOTH” CIRCUITS

Consider the case where  $A_\ell$  are product of rotations for all  $\ell$ , namely

$$A_\ell = \prod_{j=1}^P e^{-\frac{1}{2}i\beta_{\ell j}^A \theta_{\ell j}^A \mathcal{P}_{\ell j}^A}, \quad (4.1)$$

where  $\mathcal{P}_{\ell j}^A$  are Pauli matrices acting on all qubits, and similarly for  $B_\ell$ . These can also be interspersed with other non-trainable unitaries. This constitutes a slight generalization of the setting considered in Harrow & Napp (2021), and the algorithm we present is essentially a distributed distributed version of theirs. Denote by  $\beta$  an  $2PL$ -dimensional vector with elements  $\beta_{\ell j}^Q$  where  $Q \in \{A, B\}$ <sup>2</sup>. The quantity  $\|\beta\|_1$  is the total evolution time if we interpret the state  $|\varphi\rangle$  as a sequence of Hamiltonians applied to the initial state  $|x\rangle$ .

In Appendix D.1 we describe an algorithm that converges to the neighborhood of a minimum, or achieves  $\mathbb{E}\mathcal{L}(\Theta) - \mathcal{L}(\Theta^*) \leq \varepsilon_0$ , for a convex  $\mathcal{L}$  after

$$\frac{2 \|\Theta^{(0)} - \Theta^*\|_2^2 \|\beta\|_1^2}{\varepsilon_0^2} \quad (4.2)$$

---

<sup>2</sup>Harrow & Napp (2021) actually consider a related quantity for which has smaller norm in cases where multiple gradient measurements commute, leading to even better rates.

iterations, where  $\Theta^*$  are the parameter values at the minimum of  $\mathcal{L}$ . The expectation is with respect to the randomness of quantum measurement and additional internal randomness of the algorithm. The algorithm is based on classically sampling a single coordinate to update at every iteration, and computing an unbiased estimator of the gradient with a single measurement. It can thus be seen as a form of probabilistic coordinate descent.

This implies an exponential advantage in communication for the entire training process as long as  $\|\Theta^{(0)} - \Theta^*\|_2^2 \|\beta\|_1^2 = \text{polylog}(N)$ . Such circuits either have a small number of trainable parameters ( $P = O(\text{polylog}(N))$ ), else depend weakly on each parameter (e.g.  $\beta_{\ell_j}^Q = O(1/P)$  for arbitrary  $P$ ), or have structure that allows initial parameter guesses whose quality diminishes quite slowly with system size. Nevertheless, over a convex region the loss can rapidly change by an  $O(1)$  amount. One may also be concerned that in the setting  $\|\Theta^{(0)} - \Theta^*\|_2^2 \|\beta\|_1^2 = \text{polylog}(N)$  only a logarithmic number of parameters is updated during the entire training process and so the total effect of the training process may be negligible. It is important to note however that each such sparse update depends on the structure of the entire gradient vector as seen in the sampling step. In this sense the algorithm is a form of probabilistic coordinate descent, since the probability of updating a coordinate  $|\beta_{\ell_j}^Q| / \|\beta\|_1$  is proportional to the magnitude of the corresponding element in the gradient (actually serving as an upper bound for it).

Remarkably, the time complexity of a single iteration of this algorithm is proportional to a forward pass, and so matches the scaling of classical backpropagation. This is in contrast to the polynomial overhead of shadow tomography (Theorem 1). Additionally, it requires a single measurement per iteration, without any of the additional factors in the sample complexity of shadow tomography.

#### 4.2 FINE-TUNING THE LAST LAYER OF A MODEL

Consider a model given by eq. (3.1) where only the parameters of  $A_L$  are trained, and the rest are frozen, and denote this model by  $|\varphi_f\rangle$ . The circuit up to that unitary could include multiple data-dependent unitaries that represent complex features in the data. Training only the final layer in this manner is a common method of fine-tuning a pre-trained model Howard & Ruder (2018). If we now define

$$E_{L_i}^A = |1\rangle\langle 0| \otimes A_L^\dagger \mathcal{P}_0 \frac{\partial A_L}{\partial \theta_{L_i}^A} + |0\rangle\langle 1| \otimes \left( \frac{\partial A_L}{\partial \theta_{L_i}^A} \right)^\dagger \mathcal{P}_0 A_L, \quad (4.3)$$

the expectation value of  $E_{L_i}^A$  using the state  $|+\rangle |\mu^{A_L}\rangle$  gives  $\frac{\partial \mathcal{L}}{\partial \theta_{L_i}^A}$ . Here

$$|\mu^{A_L}\rangle = B_L(x) \prod_{k=L-1}^1 A_k(x) B_k(x) |\psi(x)\rangle \quad (4.4)$$

is the forward feature computed by Alice at layer  $L$  with the parameters of all the other unitaries frozen (hence the dependence on them is dropped). Since the observables in the shadow tomography problem can be chosen in an online fashion, and adaptively based on previous measurements, we can simply define a stream of measurement operators by measuring  $P$  observables to estimate the gradients w.r.t. an initial set of parameters, updating these parameters using gradient descent with step size  $\eta$ , and defining a new set of observables using the updated parameters. Repeating this for  $T$  iterations gives a total of  $PT$  observables. By the scaling in Lemma 2, the total communication needed is  $\tilde{O}(\log N (\log TP)^4 \log(L/\delta) / \varepsilon^4)$  over  $O(L)$  rounds (since only  $O(L)$  rounds are needed to create copies of  $|\mu^{A_L}\rangle$ ). This implies an exponential advantage in communication for the entire training process (under the reasonable assumption  $T = O(\text{poly}(N, P))$ ), despite the additional stochasticity introduced by the need to perform quantum measurements. For example, assume one has a bound  $\|\nabla \mathcal{L}\|_2^2 \leq K$ . If the circuit is comprised of unitaries with Hermitian derivatives, this holds with  $K = PL$ . In that case, denoting by  $g$  the gradient estimator obtained by shadow tomography, we have

$$\|g\|_2^2 \leq \|\nabla \mathcal{L}\|_2^2 + \|\nabla \mathcal{L} - g\|_2^2 \leq K + \varepsilon^2 PL. \quad (4.5)$$

It then follows directly from Lemma 6 that for an appropriately chosen step size, if  $\mathcal{L}$  is convex one can find parameter values  $\bar{\Theta}$  such that  $\mathcal{L}(\bar{\Theta}) - \mathcal{L}(\Theta^*) \leq \varepsilon_0$  using

$$T = 2 \frac{\|\Theta^{(0)} - \Theta^*\|_2^2 (K + \varepsilon^2 PL)^2}{\varepsilon_0^2} \quad (4.6)$$

iterations of gradient descent. Similarly if  $\mathcal{L}$  is  $\lambda$ -strongly convex then  $T = 2(K + \varepsilon^2 PL)^2 / \lambda \varepsilon_0 + 1$  iterations are sufficient. In both cases therefore an exponential advantage is achieved for the optimization process as a whole, since in both cases one can implement the circuit that is used to obtain the lower bounds in Lemma 3.

## 5 EXPRESSIVITY OF COMPOSITIONAL MODELS

It is natural to ask how expressive models of the form of eq. (3.1) can be, given the unitarity constraint of quantum mechanics on the matrices  $\{A_\ell, B_\ell\}$ . This is a nuanced question that can depend on the encoding of the data that is chosen and the method of readout. On the one hand, if we pick  $|\psi(x)\rangle$  as in eq. (3.5) and use  $\{A_\ell, B_\ell\}$  that are independent of  $x$ , the resulting state  $|\varphi\rangle$  will be a linear function of  $x$ . On the other hand, one could map bits to qubits 1-to-1 and encode any reversible classical function of data within the unitary matrices  $\{A_\ell(x)\}$  with the use of extra space qubits. However, this negates the possibility of any space or communication advantages (and does not provide any real computational advantage without additional processing). As above, one prefers to work on more generic functions in the amplitude and phase space, allowing for an exponential compression of the data into a quantum state, but one that must be carefully worked with.

We investigate the consequences of picking  $\{A_\ell(x)\}$  that are *nonlinear* functions of  $x$ , and  $\{B_\ell\}$  that are data-independent. This is inspired by a common use case in which Alice holds some data or features of the data, while Bob holds a model that can process these features. Given a scalar variable  $x$ , define  $A_\ell(x) = \text{diag}((e^{-2\pi i \lambda_{\ell 1} x}, \dots, e^{-2\pi i \lambda_{\ell N} x}))$  for  $\ell \in \{1, \dots, L\}$ . We also consider parameterized unitaries  $\{B_\ell\}$  that are independent of the  $\{\lambda_{\ell i}\}$  and inputs  $x, y$ , and the state obtained by interleaving the two in the manner of eq. (3.1) by  $|\varphi(x)\rangle$ .

We next set  $\lambda_{\ell 1} = 0$  for all  $\ell \in \{1, \dots, L\}$  and  $\lambda_{L 2} = 0$ . If we are interested in expressing the frequency

$$\Lambda_{\bar{j}} = \sum_{\ell=1}^{L-1} \lambda_{\ell j_\ell}, \quad (5.1)$$

where  $j_\ell \in \{2, \dots, N\}$ , we simply initialize with  $|\psi(x)\rangle = |+\rangle_0 |0\rangle$  and use

$$B_\ell = |j_\ell - 1\rangle \langle j_{\ell-1} - 1| + |j_{\ell-1} - 1\rangle \langle j_\ell - 1|, \quad (5.2)$$

with  $j_1 = j_L = 2$ . It is easy to check that the resulting state is  $|\varphi(x)\rangle = (|0\rangle + e^{-2\pi i \Lambda_{\bar{j}} x} |1\rangle) / \sqrt{2}$ . Since the basis state  $|0\rangle$  does not accumulate any phase, while the  $B_\ell$ s swap the  $|1\rangle$  state with the appropriate basis state at every layer in order to accumulate a phase corresponding to a single summand in eq. (5.1). Choosing to measure the operator  $\mathcal{P}_0 = X_0$ , it follows that  $\langle \varphi(x) | X_0 | \varphi(x) \rangle = \cos(2\pi \Lambda_{\bar{j}} x)$ .

It is possible to express  $O(N^L)$  different frequencies in this way, assuming the  $\Lambda_{\bar{j}}$  are distinct, which will be the case for example with probability 1 if the  $\{\lambda_{\ell i}\}$  are drawn i.i.d. from some distribution with continuous support. This further motivates the small  $L$  regime where exponential advantage in communication is possible. These types of circuits with interleaved data-dependent unitaries and parameterized unitaries was considered for example in Schuld et al. (2020), and is also related to the setting of quantum signal processing and related algorithms Low & Chuang (2017); Martyn et al. (2021). We also show that such circuits can express dense function in Fourier space, and for small  $N$  we additionally find that these circuits are universal function approximators (Appendix F.1), though in this setting the possible communication advantage is less clear.

The problem of applying nonlinearities to data encoded efficiently in quantum states is non-trivial and is of interest due to the importance of nonlinearities in enabling efficient function approximation Maiorov & Pinkus (1999). One approach to resolving the constraints of unitarity with the potential irreversibility of nonlinear functions is the introduction of slack variables via additional ancilla qubits, as typified by the techniques of block-encoding Chakraborty et al. (2018); Gilyén et al. (2018). Indeed, these techniques can be used to apply nonlinearities to amplitude encoded data efficiently, as was recently shown in Rattew & Reberstrost (2023). This approach can be applied to the distributed setting as well. Consider the communication problem where Alice is given  $x$  as input and Bob is given unitaries  $\{U_1, U_2\}$  over  $\log N$  qubits. Denote by  $\sigma : \mathbb{R} \rightarrow \mathbb{R}$  a nonlinear function such as the sigmoid, exponential or standard trigonometric functions, and  $n = 2^N$ . We show the following:

**Lemma 5.** *There exists a model  $|\varphi_\sigma\rangle$  of the form definition 3.1 with  $L = O(\log 1/\varepsilon)$ ,  $N' = 2^{n'}$  where  $n' = 2n + 4$  such that*

$$|\varphi_\sigma\rangle = \alpha |0\rangle^{\otimes n+4} |\hat{y}\rangle + |\phi\rangle \quad (5.3)$$



for some  $\alpha = O(1)$ , where  $|\hat{y}\rangle$  is a state that obeys

$$\left\| |\hat{y}\rangle - \left| U_2 \frac{1}{\|\sigma(U_1 x)\|_2} \sigma(U_1 x) \right\rangle \right\|_2 < \varepsilon. \quad (5.4)$$

$|\phi\rangle$  is a state whose first  $n + 4$  registers are orthogonal to  $|0\rangle^{\otimes n+4}$ .

Proof: Appendix B.

This result implies that with constant probability, after measurement of the first  $n + 4$  qubits of  $|\varphi_\sigma\rangle$ , one obtains a state whose amplitudes encode the output of a single hidden layer neural network.

It is also worth noting that the general form of the circuits we consider resembles self-attention based models with their nonlinearities removed (motivated for example by Sun et al. (2023)), as we explain in Appendix F.2. Finally, in Appendix F.3 we discuss other strategies for increasing the expressivity of these quantum circuits by combining them with classical networks.

## 6 PRIVACY OF QUANTUM COMMUNICATION

In addition to an advantage in communication complexity, the quantum algorithms outlined above have an inherent advantage in terms of privacy. It is well known that the number of bits of information that can be extracted from an unknown quantum state is proportional to the number of qubits. It follows immediately that since the above algorithm requires exchanging a logarithmic number of copies of states over  $O(\log N)$  qubits, even if all the communication between the two players is intercepted, an attacker cannot extract more than a logarithmic number of bits of classical information about the input data or model parameters. Specifically, we have:

**Corollary 1.** *If Alice and Bob are implementing the quantum algorithm for gradient estimation described in Lemma 2, and all the communication between Alice and Bob is intercepted by an attacker, the attacker cannot extract more than  $\tilde{O}(L^2(\log N)^2(\log P)^4 \log(L/\delta)/\varepsilon^4)$  bits of classical information about the inputs to the players.*

This follows directly from Holevo’s theorem Holevo (1973), since the multiple copies exchanged in each round of the protocol can be thought of as a quantum state over  $\tilde{O}((\log N)^2(\log P)^4 \log(L/\delta)/\varepsilon^4)$  qubits. As noted in Aaronson (2017b), this does not contradict the fact that the protocol allows one to estimate all  $P$  elements of the gradient, since if one were to place some distribution over the inputs, the induced distribution over the gradient elements will generally exhibit strong correlations. An analogous result holds for the inference problem described in Lemma 1.

It is also interesting to ask how much information either Bob or Alice can extract about the inputs of the other player by running the protocol. If this amount is logarithmic as well, it provides additional privacy to both the model owner and the data owner. It allows two actors who do not necessarily trust each other, or the channel through which they communicate, to cooperate in jointly training a distributed model or using one for inference while only exposing a vanishing fraction of the information they hold.

It is also worth mentioning that data privacy is also guaranteed in a scenario where the user holding the data also specifies the processing done on the data. In this setting, Alice holds both data  $x$  and a full description of the unitaries she wishes to apply to her state. She can send Bob a classical description of these unitaries, and as long as the data and features are communicated in the form of quantum states, only a logarithmic amount of information can be extracted about them. In this setting there is of course no advantage in communication complexity, since the classical description of the unitary will scale like  $\text{poly}(N, P)$ .

## 7 DISCUSSION

This work constitutes a preliminary investigation into a generic class of quantum circuits that has the potential for enabling an exponential communication advantage in problems of classical data processing including training and inference with large parameterized models over large datasets. Communication constraints may become even more relevant as models continue to grow, and in setting they are trained on data that is obtained by distributed interaction with the physical world Driess et al. (2023). Our results also naturally raise further questions regarding the expressive power and trainability of these types of circuits, which may be of independent interest. We collect some of these in Appendix H.

## REFERENCES

- Scott Aaronson. Read the fine print. *Nat. Phys.*, 11(4):291–293, April 2015.
- Scott Aaronson. Introduction to quantum information science. <https://www.scottaaronson.com/qclect.pdf>, 2017a.
- Scott Aaronson. Shadow tomography of quantum states. November 2017b.
- Amira Abbas, Robbie King, Hsin-Yuan Huang, William J Huggins, Ramis Movassagh, Dar Gilboa, and Jarrod R McClean. On quantum backpropagation, information reuse, and cheating measurement collapse. May 2023.
- Dimitris Achlioptas. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *J. Comput. System Sci.*, 66(4):671–687, June 2003.
- Alekh Agarwal, Peter L Bartlett, Pradeep Ravikumar, and Martin J Wainwright. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. September 2010.
- Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G S L Brandao, David A Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P Harrigan, Michael J Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S Humble, Sergei V Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C Platt, Chris Quintana, Eleanor G Rieffel, Pedram Roushan, Nicholas C Rubin, Daniel Sank, Kevin J Satzinger, Vadim Smelyanskiy, Kevin J Sung, Matthew D Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, October 2019.
- Koji Azuma, Sophia E Economou, David Elkouss, Paul Hilaire, Liang Jiang, Hoi-Kwong Lo, and Ilan Tzitrin. Quantum repeaters: From quantum networks to the quantum internet. December 2022.
- Krishna C Balram and Kartik Srinivasan. Piezoelectric optomechanical approaches for efficient quantum microwave-to-optical signal transduction: the need for co-design. August 2021.
- Ziv Bar-Yossef, T S Jayram, and Iordanis Kerenidis. Exponential separation of quantum and classical One-Way communication complexity. *SIAM J. Comput.*, 38(1):366–384, January 2008.
- Paul Barham, Aakanksha Chowdhery, Jeff Dean, Sanjay Ghemawat, Steven Hand, Dan Hurt, Michael Isard, Hyeontaek Lim, Ruoming Pang, Sudip Roy, Brennan Saeta, Parker Schuh, Ryan Sepassi, Laurent El Shafey, Chandramohan A Thekkath, and Yonghui Wu. Pathways: Asynchronous distributed dataflow for ML. March 2022.
- C H Bennett, G Brassard, C Crépeau, R Jozsa, A Peres, and W K Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Phys. Rev. Lett.*, 70(13):1895–1899, March 1993.
- Charles H Bennett, Gilles Brassard, Sandu Popescu, Benjamin Schumacher, John A Smolin, and William K Wootters. Purification of noisy entanglement and faithful teleportation via noisy channels. November 1995.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D Manning, Suvir

- Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W Thomas, Florian Tramèr, Rose E Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and risks of foundation models. August 2021.
- Fernando G S Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M Svore, and Xiaodi Wu. Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning. October 2017.
- Gilles Brassard. Quantum communication complexity (a survey). January 2001.
- Adam R Brown and Leonard Susskind. The second law of quantum complexity. January 2017.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are Few-Shot learners. May 2020.
- Sébastien Bubeck. Convex optimization: Algorithms and complexity. May 2014.
- Harry Buhrman, Richard Cleve, Serge Massar, and Ronald de Wolf. Non-locality and communication complexity. *arXiv [quant-ph]*, July 2009.
- M Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J Coles. Variational quantum algorithms. December 2020.
- Shantanav Chakraborty, András Gilyén, and Stacey Jeffery. The power of block-encoded matrix powers: improved regression techniques via faster hamiltonian simulation. April 2018.
- A Chi-Chih Yao. Quantum circuit complexity. In *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science*, pp. 352–361, November 1993.
- Andrew M Childs and Wim van Dam. Quantum algorithms for algebraic problems. December 2008.
- Nadav Cohen and Amnon Shashua. Inductive bias of deep convolutional networks through pooling geometry. May 2016.
- Nadav Cohen, Or Sharir, and Amnon Shashua. On the expressive power of deep learning: A tensor analysis. September 2015.
- Gavin E Crooks. Gradients of parameterized quantum gates using the parameter-shift rule and gate decomposition. May 2019.
- Danny Driess, Fei Xia, Mehdi S M Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. PaLM-E: An embodied multimodal language model. March 2023.
- Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors. February 2018.
- Richard P Feynman. Simulating physics with computers. *Int. J. Theor. Phys.*, 21(6):467–488, June 1982.
- András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. June 2018.
- Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum random access memory. *Phys. Rev. Lett.*, 100(16):160501, April 2008.

- Lukas Gonon and Antoine Jacquier. Universal approximation theorem and error bounds for quantum neural networks and quantum reservoirs. July 2023.
- Google Quantum AI. Suppressing quantum errors by scaling a surface code logical qubit. *Nature*, 614(7949):676–681, February 2023.
- Goren Gordon and Gustavo Rigolin. Generalized teleportation protocol. November 2005.
- Robert Gower, Donald Goldfarb, and Peter Richtarik. Stochastic block BFGS: Squeezing more curvature out of data. In Maria Florina Balcan and Kilian Q Weinberger (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1869–1878, New York, New York, USA, 2016. PMLR.
- Aram W Harrow and John C Napp. Low-Depth gradient measurements can improve convergence in variational hybrid Quantum-Classical algorithms. *Phys. Rev. Lett.*, 126(14):140502, April 2021.
- Aram W Harrow, Avinandan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, 103(15):150502, October 2009.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W Rae, Oriol Vinyals, and Laurent Sifre. Training Compute-Optimal large language models. March 2022.
- Alexander Semenovich Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii*, 1973.
- Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. January 2018.
- Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R McClean. Power of data in quantum machine learning. *Nat. Commun.*, 12(1):2631, May 2021.
- Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Mia Xu Chen, Dehao Chen, Hyoukjoong Lee, Jiquan Ngiam, Quoc V Le, Yonghui Wu, and Zhifeng Chen. GPipe: Efficient training of giant neural networks using pipeline parallelism. November 2018.
- William J Huggins and Jarrod R McClean. Accelerating quantum algorithms with precomputation. May 2023.
- Rahul Jain, Jaikumar Radhakrishnan, and Pranab Sen. The quantum communication complexity of the pointer chasing problem: The bit version. In *FST TCS 2002: Foundations of Software Technology and Theoretical Computer Science*, Lecture notes in computer science, pp. 218–229. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.
- Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *Lecture Notes in Computer Science*, Lecture notes in computer science, pp. 126–152. Springer International Publishing, Cham, 2018.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *NeurIPS*, 2013.
- Norman P Jouppi, George Kurian, Sheng Li, Peter Ma, Rahul Nagarajan, Lifeng Nai, Nishant Patil, Suvinay Subramanian, Andy Swing, Brian Towles, Cliff Young, Xiang Zhou, Zongwei Zhou, and David Patterson. TPU v4: An optically reconfigurable supercomputer for machine learning with hardware support for embeddings. April 2023.
- J Kaplan, S McCandlish, T Henighan, T B Brown, and others. Scaling laws for neural language models. *arXiv preprint arXiv*, 2020.
- Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are RNNs: Fast autoregressive transformers with linear attention. June 2020.
- V Krutyanskiy, M Galli, V Krcmarsky, S Baier, D A Fioretto, Y Pu, A Mazloom, P Sekatski, M Canteri, M Teller, J Schupp, J Bate, M Meraner, N Sangouard, B P Lanyon, and T E Northup. Entanglement of trapped-ion qubits separated by 230 meters. August 2022.

- Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, Cambridge, England, April 2011.
- Nikolai Lauk, Neil Sinclair, Shabir Barzanjeh, Jacob P Covey, Mark Saffman, Maria Spiropulu, and Christoph Simon. Perspectives on quantum transduction. *Quantum Sci. Technol.*, 5(2):020501, March 2020.
- Nicolas Le Roux, Mark Schmidt, and Francis Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. February 2012.
- James Lee-Thorp, Joshua Ainslie, Ilya Eckstein, and Santiago Ontanon. FNet: Mixing tokens with fourier transforms. May 2021.
- Yoav Levine, Or Sharir, Alon Ziv, and Amnon Shashua. On the Long-Term memory of deep recurrent networks. October 2017.
- Yoav Levine, Noam Wies, Or Sharir, Hofit Bata, and Amnon Shashua. The Depth-to-Width interplay in Self-Attention. June 2020.
- Bo Li, Yuan Cao, Yu-Huai Li, Wen-Qi Cai, Wei-Yue Liu, Ji-Gang Ren, Sheng-Kai Liao, Hui-Nan Wu, Shuang-Lin Li, Li Li, Nai-Le Liu, Chao-Yang Lu, Juan Yin, Yu-Ao Chen, Cheng-Zhi Peng, and Jian-Wei Pan. Quantum state transfer over 1200 km assisted by prior distributed entanglement. *Phys. Rev. Lett.*, 128(17):170501, April 2022.
- S Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, August 1996.
- Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nat. Phys.*, 10(9):631–633, July 2014.
- Guang Hao Low and Isaac L Chuang. Optimal hamiltonian simulation by quantum signal processing. *Phys. Rev. Lett.*, 118(1):010501, January 2017.
- P Magnard, S Storz, P Kurpiers, J Schär, F Marxer, J Lütolf, T Walter, J-C Besse, M Gabureac, K Reuer, A Akin, B Royer, A Blais, and A Wallraff. Microwave quantum link between superconducting circuits housed in spatially separated cryogenic systems. *Phys. Rev. Lett.*, 125(26):260502, December 2020.
- Vitaly Maiorov and Allan Pinkus. Lower bounds for approximation by MLP neural networks. *Neurocomputing*, 25(1):81–91, April 1999.
- John M Martyn, Zane M Rossi, Andrew K Tan, and Isaac L Chuang. A grand unification of quantum algorithms. May 2021.
- Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. September 2015.
- Hagay Michaeli, Tomer Michaeli, and Daniel Soudry. Alias-Free convnets: Fractional shift invariance via polynomial activations. March 2023.
- Boris Mityagin. The zero set of a real analytic function. December 2015.
- Ashley Montanaro and Sam Pallister. Quantum algorithms and the finite element method. December 2015.
- Philipp Moritz, Robert Nishihara, and Michael Jordan. A Linearly-Convergent stochastic L-BFGS algorithm. In Arthur Gretton and Christian C Robert (eds.), *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pp. 249–258, Cadiz, Spain, 2016. PMLR.
- Danial Motlagh and Nathan Wiebe. Generalized quantum signal processing. August 2023.
- William J Munro, Koji Azuma, Kiyoshi Tamaki, and Kae Nemoto. Inside quantum repeaters. *IEEE J. Sel. Top. Quantum Electron.*, 21(3):78–90, May 2015.

- Deepak Narayanan, Aaron Harlap, Amar Phanishayee, Vivek Seshadri, Nikhil R Devanur, Gregory R Ganger, Phillip B Gibbons, and Matei Zaharia. PipeDream: generalized pipeline parallelism for DNN training. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP '19*, pp. 1–15, New York, NY, USA, October 2019. Association for Computing Machinery.
- Ashwin Nayak and Felix Wu. The quantum query complexity of approximating the median and related statistics. April 1998.
- Michael A Nielsen and Isaac L Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, December 2010.
- Brad G Osgood. *Lectures on the Fourier Transform and Its Applications (Pure and Applied Undergraduate Texts) (Pure and Applied Undergraduate Texts, 33)*. American Mathematical Society, January 2019.
- Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I Latorre. Data re-uploading for a universal quantum classifier. July 2019.
- M Pompili, S L N Hermans, S Baier, H K C Beukers, P C Humphreys, R N Schouten, R F L Vermeulen, M J Tiggelman, L Dos Santos Martins, B Dirkse, S Wehner, and R Hanson. Realization of a multinode quantum network of remote solid-state qubits. *Science*, 372(6539):259–264, April 2021.
- Stephen J Ponzio, Jaikumar Radhakrishnan, and S Venkatesh. The communication complexity of pointer chasing. *J. Comput. System Sci.*, 62(2):323–355, March 2001.
- Reiner Pope, Sholto Douglas, Aakanksha Chowdhery, Jacob Devlin, James Bradbury, Anselm Levskaya, Jonathan Heek, Kefan Xiao, Shivani Agrawal, and Jeff Dean. Efficiently scaling transformer inference. November 2022.
- Anup Rao and Amir Yehudayoff. *Communication Complexity and Applications*. Cambridge University Press, January 2020.
- Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. DeepSpeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20*, pp. 3505–3506, New York, NY, USA, August 2020. Association for Computing Machinery.
- Arthur G Rattew and Patrick Rebentrost. Non-Linear transformations of quantum amplitudes: Exponential improvement, generalization, and applications. September 2023.
- Ran Raz. Exponential separation of quantum and classical communication complexity. In *Proceedings of the thirty-first annual ACM symposium on Theory of Computing, STOC '99*, pp. 358–367, New York, NY, USA, May 1999. Association for Computing Machinery.
- Alexander Razborov. Quantum communication complexity of symmetric predicates. April 2002.
- Tim Roughgarden. Communication complexity (for algorithm designers). September 2015.
- Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. The effect of data encoding on the expressive power of variational quantum machine learning models. August 2020.
- P W Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*. IEEE Comput. Soc. Press, 1994.
- Yutao Sun, Li Dong, Shaohan Huang, Shuming Ma, Yuqing Xia, Jilong Xue, Jianyong Wang, and Furu Wei. Retentive network: A successor to transformer for large language models. July 2023.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł Ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, and R Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30, pp. 5998–6008. Curran Associates, Inc., 2017.
- A J Walker. New fast method for generating discrete random numbers with arbitrary frequency distributions. *Electron. Lett.*, 10(8):127–128, April 1974.

Changqing Wang, Ivan Gonin, Anna Grassellino, Sergey Kazakov, Alexander Romanenko, Vyacheslav P Yakovlev, and Silvia Zorzetti. High-efficiency microwave-optical quantum transduction based on a cavity electro-optic superconducting system with long coherence time. *npj Quantum Information*, 8(1):1–10, December 2022.

Yuanzhong Xu, Hyoukjoong Lee, Dehao Chen, Blake Hechtman, Yanping Huang, Rahul Joshi, Maxim Krikun, Dmitry Lepikhin, Andy Ly, Marcello Maggioni, Ruoming Pang, Noam Shazeer, Shibo Wang, Tao Wang, Yonghui Wu, and Zhifeng Chen. GSPMD: General and scalable parallelization for ML computation graphs. May 2021.

Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the eleventh annual ACM symposium on Theory of computing, STOC '79*, pp. 209–213, New York, NY, USA, April 1979. Association for Computing Machinery.

## A NOTATION AND REVIEW OF QUANTUM MECHANICS

We denote by  $\{a_i\}$  a set of elements indexed by  $i$ , with 1-based indexing unless otherwise specified, with the maximal value of  $i$  explicitly specified when it is not clear from context.  $[N]$  denotes the set  $\{0, \dots, N-1\}$ . The complex conjugate of a number  $c$  is denoted by  $c^*$ , and the conjugate transpose of a complex-valued matrix  $A$  by  $A^\dagger$ .

We denote by  $|\psi\rangle$  a vector of complex numbers  $\{\psi_i\}$  representing the state of a quantum system when properly normalized, and by  $\langle\psi|$  its dual (assuming it exists). The inner product between two such vectors of length  $N$  is denoted by

$$\langle\psi|\varphi\rangle = \sum_{i=0}^{N-1} \psi_i^* \varphi_i. \quad (\text{A.1})$$

Denoting by  $|i\rangle$  for  $i \in [N]$  a basis vector in an orthonormal basis with respect to the above inner product, we can also write

$$|\psi\rangle = \sum_{i=0}^{N-1} \psi_i |i\rangle. \quad (\text{A.2})$$

Matrices will be denoted by capital letters, and when acting on quantum states will always be unitary. These can be specified in terms of their matrix elements using the Dirac notation defined above, as in

$$A = \sum_{ij} A_{ij} |i\rangle \langle j|. \quad (\text{A.3})$$

Matrix-vector product are specified naturally in this notation by

Quantum mechanics is, in the simplest possible terms, a theory of probability based on conservation of the  $L^2$  norm rather than the standard probability theory based on the  $L^1$  norm Aaronson (2017a); Nielsen & Chuang (2010). The state of a pure quantum system is described fully by a complex vector of  $N$  numbers known as amplitudes which we denote by  $\{\psi_i\}$  where  $i \in \{0, \dots, N-1\}$ , and is written using Dirac notation as  $|\psi\rangle$ . The state is normalized so that

$$\langle\psi|\psi\rangle = \sum_{i=0}^{N-1} \psi_i^* \psi_i = \sum_{i=0}^{N-1} |\psi_i|^2 = 1, \quad (\text{A.4})$$

which is the  $L^2$  equivalent of the standard normalization condition of classical probability theory. It is a curious fact that the choice of  $L^2$  requires the use of complex rather than real amplitudes, and that no consistent theory can be written in this way for any other  $L^p$  norm Aaronson (2017a). The most general state of a quantum system is a probabilistic mixture of pure states, in the sense of the standard  $L^1$ -based rules of probability. We will not be concerned with these types of states, and so omit their description here, and subsequently whenever quantum states are discussed, the assumption is that they are pure.

Since any closed quantum system conserves probability, the  $L^2$  norm of a quantum state is conserved during the evolution of a quantum state. Consequently, when representing and manipulating quantum states on a quantum computer, the fundamental operation is the application of a unitary matrix to a quantum state.

Given a quantum system with some discrete degrees of freedom, the number of amplitudes corresponds to the number of possible states of the system, and is thus exponential in the number of degrees of freedom. The simplest such degree of freedom is a binary one, called a qubit, which is analogous to a bit. Thus a state of  $\log N$  qubits is described by  $N$  complex amplitudes.

A fundamental property of quantum mechanics is that the amplitudes of a quantum state are not directly measurable. Given a Hermitian operator

$$\mathcal{O} = \sum_{i=0}^{N-1} \lambda_i |v_i\rangle \langle v_i| \quad (\text{A.5})$$

with real eigenvalues  $\{\lambda_i\}$ , a measurement of  $\mathcal{O}$  with respect to a state  $|\psi\rangle$  gives the result  $\lambda_i$  with probability  $|\langle v_i|\psi\rangle|^2$ . The real-valued quantity

$$\langle\psi|\mathcal{O}|\psi\rangle = \sum_{i=0}^{N-1} \lambda_i |\langle\psi|v_i\rangle|^2 \quad (\text{A.6})$$



is the expectation value of  $\mathcal{O}$  with respect to  $|\psi\rangle$ , and its value can be estimated by measurements. After a measurement with outcome  $\lambda_i$ , the original state is destroyed, collapsing to the state  $|v_i\rangle$ . This fundamentally destructive nature of quantum measurement implies that simply encoding information in the amplitudes of a quantum state does not necessarily render it useful for downstream computation. The design of quantum algorithms essentially amounts to a careful and intricate design of amplitude manipulations and measurements in order to extract useful information from the amplitudes of a quantum state.

## B PROOFS

*Proof of Lemma 1.* Clearly  $\langle\varphi|\mathcal{P}_0|\varphi\rangle$  can be estimated by preparing  $|\varphi\rangle$  and measuring it  $O(1/\varepsilon^2)$  times. Preparing each copy of  $|\varphi\rangle$  requires  $O(L)$  rounds of communication, with each round involving the communication of a  $\log N'$ -qubit quantum state. Alice first prepares  $|\psi(x)\rangle$ , and this state is passed back and forth with each player applying  $A_\ell$  or  $B_\ell$  respectively for  $\ell \in \{1, \dots, L\}$ . □

*Proof of Lemma 2.* The proof reduces gradient computation the problem of estimating the expectation values of a set of observables, as was done e.g. in Abbas et al. (2023); Harrow & Napp (2021). We consider the parameters of the unitaries that Alice possesses first, and an identical argument follows for the parameters of Bob's unitaries.

We have

$$\begin{aligned} \frac{\partial}{\partial\theta_{\ell i}^A} \langle\varphi|\mathcal{P}_0|\varphi\rangle &= 2\text{Im} \langle\varphi|\mathcal{P}_0 \prod_{k=L}^{\ell+1} A_k B_k \frac{\partial A_\ell}{\partial\theta_{\ell i}^A} B_\ell \prod_{k=\ell-1}^1 A_k B_k |\psi(x)\rangle \\ &\equiv 2\text{Im} \langle\nu_{\ell i}^A|\mu_\ell^A\rangle \end{aligned} \quad (\text{B.1})$$

where

$$|\mu_\ell^A\rangle = B_\ell \prod_{k=\ell-1}^1 A_k B_k |\psi(x)\rangle, \quad |\nu_{\ell i}^A\rangle = \left(\frac{\partial A_\ell}{\partial\theta_{\ell i}^A}\right)^\dagger \prod_{k=L}^{\ell+1} B_k^\dagger A_k^\dagger \mathcal{P}_0 |\varphi\rangle, \quad (\text{B.2})$$

correspond to forward and backward features for the  $i$ -th parameter of  $A_\ell$  respectively. This is illustrated graphically in Figure 1. We also write

$$|\nu_{\ell 0}^A\rangle = \prod_{k=L}^{\ell+1} B_k^\dagger A_k^\dagger \mathcal{P}_0 |\varphi\rangle. \quad (\text{B.3})$$

Attaching an ancilla qubit denoted by  $a$  to the feature states defined above, we define

$$|\psi_{\ell i}^A\rangle \equiv \frac{1}{\sqrt{2}} (|0\rangle |\mu_\ell^A\rangle + |1\rangle |\nu_{\ell i}^A\rangle), \quad (\text{B.4})$$

and a Hermitian measurement operator

$$\begin{aligned} E_{\ell i}^A &\equiv \left(|0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes \left(\frac{\partial A_\ell}{\partial\theta_{\ell i}^A}\right)\right) X_a \left(|0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes \left(\frac{\partial A_\ell}{\partial\theta_{\ell i}^A}\right)^\dagger\right) \\ &= |1\rangle\langle 0| \otimes \left(\frac{\partial A_\ell}{\partial\theta_{\ell i}^A}\right) + |0\rangle\langle 1| \otimes \left(\frac{\partial A_\ell}{\partial\theta_{\ell i}^A}\right)^\dagger, \end{aligned} \quad (\text{B.5})$$

we then have

$$\begin{aligned} \langle\psi_{\ell 0}^A|E_{\ell i}^A|\psi_{\ell 0}^A\rangle &= \langle\psi_{\ell i}^A|X_a|\psi_{\ell i}^A\rangle \\ &= \frac{\partial}{\partial\theta_{\ell i}^A} \langle\varphi|\mathcal{P}_0|\varphi\rangle, \end{aligned} \quad (\text{B.6})$$

where  $X_a$  acts on the ancilla.

Note that  $|\psi_{\ell 0}^A\rangle^{\otimes k}$  can be prepared by Alice first preparing  $(|+\rangle|\psi(x)\rangle)^{\otimes k}$  and sending this state back and forth at most  $2L$  times, with each player applying the appropriate unitaries conditioned on the value of the ancilla. Additionally,

for any choice of  $\ell$  and any  $i$ , Alice has full knowledge of the  $E_{\ell i}^A$ . They can thus be applied to quantum states and classical hypothesis states without requiring any communication.

The gradient can then be estimated using shadow tomography (Theorem 1). Specifically, for each  $\ell$ , Alice prepares  $\tilde{O}(\log^4 P \log N \log(L/\delta)/\varepsilon^4)$  copies of  $|\psi_0^{A_\ell}\rangle$ , which requires  $O(L)$  rounds of communication, each of  $\tilde{O}(\log^4 P \log^2 N \log(L/\delta)/\varepsilon^4)$  qubits. She then runs shadow tomography to estimate  $\nabla_{A_\ell} \langle \varphi | Z_0 | \varphi \rangle$  up to error  $\varepsilon$  with no additional communication. Bob does the same to estimate  $\nabla_{B_\ell} \langle \varphi | Z_0 | \varphi \rangle$ . In total  $O(L^2)$  rounds are needed to estimate the full gradient. The success probability of all  $L$  applications of shadow tomography is at least  $1 - \delta$  by a union bound.

Based on the results of Brandão et al. (2017), the space and time complexity of each application of shadow tomography is  $\sqrt{P} \text{poly}(N, \log P, \varepsilon^{-1}, \log(1/\delta))$ . This is the query complexity of the algorithm to oracles that implement the measurement operators  $\{E_{\ell i}^Q\}$ . Instantiating these oracles will incur a cost of at most  $O(N^2)$ . In cases where these operators have low rank the query complexity will depend polynomially only on the rank instead of on  $N$ .  $\square$

*Proof of Lemma 3.* We first prove an  $\Omega(\sqrt{N})$  lower bound on the amount of classical communication. Consider the following problem:

**Problem 3** (Raz (1999)). *Alice is given a vector  $x \in S^{N-1}$  and two orthogonal linear subspaces of  $\mathbb{R}^N$  each of dimension  $N/2$ , denoted  $M_1, M_2$ . Bob is given an orthogonal matrix  $O$ . Alice and Bob must determine whether  $Ox \in M_1$  or  $Ox \in M_2$ , under a promise that  $Ox$  has large support on one of the two subspaces.*

Raz (1999) showed that the randomized<sup>3</sup> classical communication complexity of the problem is  $\Omega(\sqrt{N})$ . We henceforth disregard the promise, since doing this makes the problem harder and cannot invalidate the lower bound.

The reduction from Problem 3 to Problem 1 is obtained by simply setting  $L = 1, B_1 = O, |\psi(x)\rangle = |x\rangle, \mathcal{P}_0 = Z_0$ , and

$$A_1 = \sum_{j=0}^{N/2-1} |0\rangle |j\rangle \langle v_j^1| + \sum_{j=0}^{N/2-1} |1\rangle |j\rangle \langle v_j^2|, \quad (\text{B.7})$$

where the first register contains a single qubit and  $\{v_j^k\}$  form an orthonormal basis of  $M_k$ , and picking any  $\varepsilon < 1$ . Estimating  $\mathcal{L}$  to this accuracy now solves the desired problem since  $\mathcal{L} = \langle x | O^T (\Pi_1 - \Pi_2) O | x \rangle$  where  $\Pi_k$  is a projector onto  $M_k$ , and hence determining its sign allows Alice and Bob to determine which subspace has large overlap with  $Ox$ .

The reduction from Problem 3 to Problem 2 is obtained by setting  $L = 2$ , picking  $A_1, B_1$  as before, and additionally  $B_2 = I, A_2 = e^{-i\theta_{2,1}^A X_0/2}$  initialized at  $\theta_{2,1}^A = -\pi/2$ . By the parameter shift rule Crooks (2019), we have that if  $U = e^{-i\theta \mathcal{P}}$  for some Pauli matrix  $\mathcal{P}$ , and  $U$  is part of the parameterized circuit that defines  $|\varphi\rangle$ , then

$$\frac{\partial \mathcal{L}}{\partial \theta} = \frac{1}{2} (\mathcal{L}(\theta + \frac{\pi}{2}) - \mathcal{L}(\theta - \frac{\pi}{2})). \quad (\text{B.8})$$

It follows that

$$\begin{aligned} \left. \frac{\partial \mathcal{L}}{\partial \theta_{2,1}^A} \right|_{\theta_{2,1}^A = -\pi/2} &= \frac{1}{2} (\mathcal{L}(0) - \mathcal{L}(-\pi)) \\ &= \frac{1}{2} \left( \mathcal{L}(0) - \langle x | B_1^\dagger A_1^\dagger e^{-i\frac{\pi}{2} X_0} Z_0 e^{i\frac{\pi}{2} X_0} A_1 B_1 | x \rangle \right) \\ &= \frac{1}{2} \left( \mathcal{L}(0) - \langle x | B_1^\dagger A_1^\dagger X_0 Z_0 X_0 A_1 B_1 | x \rangle \right) \\ &= \frac{1}{2} \left( \mathcal{L}(0) + \langle x | B_1^\dagger A_1^\dagger Z_0 A_1 B_1 | x \rangle \right) \\ &= \mathcal{L}(0). \end{aligned} \quad (\text{B.9})$$

<sup>3</sup>In this setting Alice and Bob can share an arbitrary number of random bits that are independent of their inputs.

Estimating  $\nabla_A \langle \varphi | Z_0 | \varphi \rangle$  to accuracy  $\varepsilon < 2$  allows one to determine the sign of  $\mathcal{L}(0)$ , which as before gives the solution to Problem 3.

Next, we show that  $\Omega(L)$  rounds are necessary in both the quantum and classical setting by a reduction from the bit version of pointer-chasing, as studied in Jain et al. (2002); Ponzio et al. (2001).

**Problem 4** (Pointer-chasing, bit version). *Alice receives a function  $f_A : [N] \rightarrow [N]$  and Bob receives a function  $f_B : [N] \rightarrow [N]$ . Alice is also given a starting point  $x \in [N]$ , and both receive an integer  $L_0$ . Their goal is to compute the least significant bit of  $f^{(L_0)}(x)$ , where  $f^{(1)}(x) = f_B(x)$ ,  $f^{(2)}(x) = f_A(f_B(x))$ ,  $\dots$*

Jain et al. (2002) show that the quantum communication complexity of  $L_0$ -round bit pointer-chasing when Bob speaks first is  $\Omega(N/L_0^4)$  (which holds for classical communication as well). This also bounds the  $L_0 - 1$ -round complexity when Alice speaks first (since such a protocol is strictly less powerful given that there are fewer rounds of communication). On the other hand, there is a trivial  $L_0$ -round protocol when Alice speaks first that requires  $\log N$  bits of communication per round, in which Alice sends Bob  $x$ , he sends back  $f^{(1)}(x)$ , she replies with  $f^{(2)}(f^{(1)}(x))$ , and so forth. This, combined with the lower bound, implies an exponential separation in communication complexity as a function of the number of rounds.

To reduce this problem to Problem 1, we assume  $f_A, f_B$  are invertible. This should not make the problem any easier since it implies that  $f_A, f_B$  have the largest possible image. In this setting,  $f_A, f_B$  can be described by unitary permutation matrices:

$$U_A = \sum_i |f_A(i)\rangle \langle i|, U_B = \sum_i |f_B(i)\rangle \langle i|. \quad (\text{B.10})$$

The corresponding circuit eq. (3.6) is then given by

$$|\varphi\rangle = \text{SWAP}_{0 \leftrightarrow \log N - 1} U_B \dots U_A U_B |x\rangle \quad (\text{B.11})$$

in the case where Bob applies the function last, with an analogous circuit in the converse situation (if Bob performed the swap, Alice applies an additional identity map). Estimating  $Z_0$  to accuracy  $\varepsilon < 1$  using this state will then reveal the least significant bit of  $f^{(L_0)}(x)$ . This gives a circuit with  $L$  layers, where  $L_0 \leq 2L - 1$ . Thus any protocol with less than  $L_0$  rounds (meaning less than  $2L - 1$  rounds) would require communicating  $\Omega(N/L_0^4) = \Omega(N/L^4)$  qubits, since the converse will contradict the results of Jain et al. (2002). The reduction to Problem 2 is along similar lines to the one described by eq. (B.9), with the state in that circuit replaced by eq. (B.11). This requires at most two additional rounds of communication.

Since quantum communication is at least as powerful than classical communication, these bounds also hold for classical communication. Since each round involves communicating at least a single bit, this gives an  $\Omega(L)$  bound on the classical communication complexity.  $\square$

*Proof of Lemma 8.* Consider first a single variable  $z$ , with data-dependent unitaries given by eq. (F.1a). If  $\{\lambda_{\ell i}\}$  are chosen i.i.d. from a uniform distribution over say  $[0, 1]$ , then with probability 1 they are all unique and so are all sums of the form  $\Lambda_{\bar{j}} = \sum_{\ell=1}^L \lambda_{\ell j_\ell}$  as well as differences  $\Lambda_{\bar{j}} - \Lambda_{\bar{k}}$  for  $\bar{k} < \bar{j}$  where the inequality holds element-wise. Set  $B_\ell$  to be the Hadamard transform over  $\log N$  qubits for all  $\ell$ , and pick the measurement operator  $\mathcal{P}_0 = X_0$ . We then have

$$\begin{aligned} \mathcal{L}_1 &= \langle \varphi | X_0 | \varphi \rangle \\ &= \sum_{\bar{j}, \bar{k} \in [N]^L} e^{2\pi i (\Lambda_{\bar{j}} - \Lambda_{\bar{k}}) z} \left( B_1^\dagger \right)_{1j_1} \left( B_2^\dagger \right)_{j_1 j_2} \dots \left( B_L^\dagger \right)_{j_{L-1} j_L} (X_0)_{j_L k_L} (B_L)_{k_L k_{L-1}} \dots (B_1)_{k_1 1} \\ &= \sum_{\bar{j}, \bar{k} \in [N]^L, \bar{j} \neq \bar{k}} e^{2\pi i (\Lambda_{\bar{j}} - \Lambda_{\bar{k}}) z} \left( B_1^\dagger \right)_{1j_1} \left( B_2^\dagger \right)_{j_1 j_2} \dots \left( B_L^\dagger \right)_{j_{L-1} j_L} (X_0)_{j_L k_L} (B_L)_{k_L k_{L-1}} \dots (B_1)_{k_1 1} \\ &= \sum_{\bar{j} \in [N]^L} \sum_{\bar{k} < \bar{j}} 2 \cos \left( 2\pi (\Lambda_{\bar{j}} - \Lambda_{\bar{k}}) z \right) \left( B_1^\dagger \right)_{1j_1} \dots \left( B_L^\dagger \right)_{j_{L-1} j_L} (X_0)_{j_L k_L} \dots (B_1)_{k_1 1} \\ &= \sum_{\bar{j}[-1] \in [N]^{L-1}} \sum_{\bar{k}[-1] < \bar{j}[-1]} \sum_{j_L=1}^N 2 \cos \left( 2\pi (\Lambda_{\bar{j}[-1]} - \Lambda_{\bar{k}[-1]} + \lambda_{L j_L} - \lambda_{L \bar{j}_L}) z \right) \\ &\quad * \left( B_1^\dagger \right)_{1j_1} \dots \left( B_L^\dagger \right)_{j_{L-1} j_L} (B_L)_{\bar{j}_L, k_{L-1}} \dots (B_1)_{k_1 1} \end{aligned} \quad (\text{B.12})$$

where  $\tilde{j}_L = j_L + (-1)^{\lfloor j_L/(N/2+1) \rfloor} N/2$ . In the third line, we dropped the diagonal terms in the double sum since they vanish due to the  $X_0$  matrix having 0 on its diagonal. In the fourth line, we collected terms and used the symmetry of  $(B_1^\dagger)_{1j_1} \dots (B_L^\dagger)_{j_{L-1}j_L} (X_0)_{j_L k_L} \dots (B_1)_{k_1 1}$  to the permutation of  $\bar{j}$  and  $\bar{k}$ . In the last line we performed the sum over  $k_L$  using the structure of  $X_0$ . By our assumption about the  $\{\lambda_{\ell_i}\}$ , each term in the final sum has a unique frequency so no cancellations are possible. The coefficient of each cosine is nonzero (and is equal to  $2N^{-L}$  or  $-2N^{-L}$ ). There are a total of  $\left(\frac{N(N-1)}{2}\right)^{L-1} N$  such summands. This completes the first part of the proof for this choice of  $\{B_\ell\}$ .

Considering instead the case of two variables, with unitaries given by eq. (F.1b), an equivalent calculation gives

$$\mathcal{L}_2 = \sum_{\bar{j}[: -1] \in [N]^{L-1}} \sum_{\bar{k}[: -1] < \bar{j}[: -1]} \sum_{j_L=1}^N 2 \cos\left(\omega_{\bar{j}\bar{k}}^1 y + \omega_{\bar{j}\bar{k}}^2 z\right) (B_1^\dagger)_{1j_1} \dots (B_L^\dagger)_{j_{L-1}j_L} (B_L)_{\bar{j}_L, k_{L-1}} \dots (B_1)_{k_1 1}, \quad (\text{B.13})$$

where

$$\omega_{\bar{j}\bar{k}}^1 = 2\pi \left( \Lambda_{\bar{j}[: N/2+1]} - \Lambda_{\bar{k}[: N/2+1]} \right), \quad \omega_{\bar{j}\bar{k}}^2 = 2\pi \left( \Lambda_{\bar{j}[N/2+1: -1]} - \Lambda_{\bar{k}[N/2+1: -1]} + \lambda_{Lj_L} - \lambda_{L\bar{j}_L} \right). \quad (\text{B.14})$$

As before, there are  $\left(\frac{N(N-1)}{2}\right)^{L-1} N$  summands in total. Since

$$\cos\left(\omega_{\bar{j}\bar{k}}^1 y + \omega_{\bar{j}\bar{k}}^2 z\right) = \cos\left(\omega_{\bar{j}\bar{k}}^1 y\right) \cos\left(\omega_{\bar{j}\bar{k}}^2 z\right) - \sin\left(\omega_{\bar{j}\bar{k}}^1 y\right) \sin\left(\omega_{\bar{j}\bar{k}}^2 z\right), \quad (\text{B.15})$$

we can rewrite eq. (B.13) as a sum over  $2\left(\frac{N(N-1)}{2}\right)^{L-1} N$  terms that are pairwise orthogonal w.r.t. the  $L^2$  inner product over  $\mathbb{R}^2$ . It follows from the definition of the separation rank that

$$\text{sep}(\mathcal{L}_2; y, z) = 2 \left( \frac{N(N-1)}{2} \right)^{L-1} N. \quad (\text{B.16})$$

We next use the assumption that the real and imaginary parts of each element of  $B_\ell$  are real analytic function of parameters  $\Theta$ . This implies that the same property holds for product of entries of the form

$$(B_1^\dagger)_{1j_1} \dots (B_L^\dagger)_{j_{L-1}j_L} (B_L)_{\bar{j}_L, k_{L-1}} \dots (B_1)_{k_1 1} \quad (\text{B.17})$$

for any choice of  $\bar{j}, \bar{k}$ . This coefficient is equal to 0 iff both the real and imaginary parts are equal to 0. Since the zero set of a real analytic function has measure 0 Mityagin (2015), the set of values of  $\Theta$  for which any of the coefficients in eq. (B.13) vanishes also has measure 0, for all choices of  $\bar{j}, \bar{k}$ . The result follows.  $\square$

*Proof of Lemma 9.* Consider a periodic function  $f$  with period 1. Denote by  $S_M[f]$  the truncated Fourier series of  $f$  written in terms of trigonometric functions:

$$\begin{aligned} S_M[f](y) &= \sum_{m=0}^{M-1} \int_{x=-1/2}^{1/2} f(x) \cos(2\pi m x) dx \cos(2\pi m y) + \sum_{m=0}^{M-1} \int_{x=-1/2}^{1/2} f(x) \sin(2\pi m x) dx \sin(2\pi m y) \\ &\equiv \sum_{m=0}^{M-1} \hat{f}_m^+ \cos(2\pi m y) + \sum_{m=0}^{M-1} \hat{f}_m^- \sin(2\pi m y). \end{aligned} \quad (\text{B.18})$$

If  $f$  is  $p$ -times continuously differentiable, it is known that the Fourier series converges uniformly, with rate

$$\|S_M[f] - f\|_\infty < \frac{C}{M^{p-1/2}}. \quad (\text{B.19})$$

for some absolute constant  $C$  Osgood (2019). For analytic functions the rate is exponential in  $M$ .

We now define the following circuit:

$$A_1(x) = \text{diag}\left(\underbrace{(1, \dots, 1)}_{N/2}, \underbrace{1, e^{2\pi i x}, e^{2\pi i 2x}, \dots, e^{2\pi i(N/4-1)x}}_{N/4}, \underbrace{1, e^{2\pi i x}, e^{2\pi i 2x}, \dots, e^{2\pi i(N/4-1)x}}_{N/4}\right), \quad (\text{B.20})$$

$$B_1 = |\hat{f}\rangle \langle 0| + |0\rangle \langle \hat{f}|, \quad (\text{B.21})$$

where

$$|\hat{f}\rangle = \frac{1}{\sqrt{\sum_m |\hat{f}_m|}} \sum_{m=0}^{N/4-1} \left( \sqrt{\hat{f}_m^+} \frac{|0\rangle + \text{sign}(\hat{f}_m^+) |1\rangle}{\sqrt{2}} |0\rangle + \sqrt{\hat{f}_m^-} \frac{|0\rangle - i \text{sign}(\hat{f}_m^-) |1\rangle}{\sqrt{2}} |1\rangle \right) |m\rangle. \quad (\text{B.22})$$

Choosing  $|\psi(x)\rangle = |0\rangle$  as the initial state, this gives

$$\begin{aligned} |\varphi\rangle &= A_1 B_1 |0\rangle \\ &= A_1 |\hat{f}\rangle \\ &= \frac{1}{\sqrt{\sum_m |\hat{f}_m|}} \sum_{m=0}^{N/4-1} \left( \sqrt{\hat{f}_m^+} \frac{|0\rangle + \text{sign}(\hat{f}_m^+) e^{2\pi i m x} |1\rangle}{\sqrt{2}} |0\rangle + \sqrt{\hat{f}_m^-} \frac{|0\rangle - i \text{sign}(\hat{f}_m^-) e^{2\pi i m x} |1\rangle}{\sqrt{2}} |1\rangle \right) |m\rangle \end{aligned} \quad (\text{B.23})$$

It follows that

$$\begin{aligned} \langle \varphi | X_0 | \varphi \rangle &= \frac{1}{\sum_m |\hat{f}_m|} \sum_{m=0}^{N/4-1} \left[ |\hat{f}_m^+| \frac{\langle 0| + \text{sign}(\hat{f}_m^+) e^{-2\pi i m x} \langle 1|}{\sqrt{2}} X_0 \frac{|0\rangle + \text{sign}(\hat{f}_m^+) e^{2\pi i m x} |1\rangle}{\sqrt{2}} \right. \\ &\quad \left. + |\hat{f}_m^-| \frac{\langle 0| + i \text{sign}(\hat{f}_m^-) e^{-2\pi i m x} \langle 1|}{\sqrt{2}} X_0 \frac{|0\rangle - i \text{sign}(\hat{f}_m^-) e^{2\pi i m x} |1\rangle}{\sqrt{2}} \right] \\ &= \frac{1}{\sum_m |\hat{f}_m|} \sum_{m=0}^{N/4-1} \hat{f}_m^+ \cos(2\pi m x) + \hat{f}_m^- \sin(2\pi m x) \\ &= \frac{1}{\sum_m |\hat{f}_m|} S_{N/4}[f](x) \end{aligned} \quad (\text{B.24})$$

This approximation thus converges uniformly according to eq. (B.19), with error decaying exponentially with number of qubits  $\log N$  as long as  $f$  is continuously differentiable at least once.  $\square$

*Proof of Lemma 5.* The algorithm in Theorem 5 of Rattew & Reberst (2023) takes as input a state-preparation unitary  $U$  acting on  $n = \log N$  qubits such that  $U |0\rangle^{\otimes n} = |z\rangle$ . Using  $O(\log 1/\varepsilon)$  queries to  $U$  and  $U^\dagger$  and  $n + 4$  ancillas, it creates a state  $|\varphi\rangle$  such that measuring 0 on the first  $n + 4$  qubits of  $|\varphi\rangle$  results in a state  $|\hat{\varphi}\rangle$  that obeys

$$\left\| |\hat{\varphi}\rangle - \frac{1}{\|\sigma(z)\|_2} |\sigma(z)\rangle \right\|_2 < \varepsilon. \quad (\text{B.25})$$

Additionally, the probability of measuring 0 on the first  $n + 4$  qubits is  $O(1)$ .

We will be interested in applying this algorithm to the state  $|U_1 x\rangle$ . The state preparation unitary can be instantiated with a single round of communication by Alice starting with the state  $|0\rangle^{\otimes 2n+4}$ , applying a unitary that encodes  $x$  in the last  $n$  qubits of this state, and then sending it to Bob who applies  $U_1$  to the same  $n$  qubits. The conjugate of the state-preparation unitary can be applied in a similar fashion by reversing this procedure. This can include any conditioning required on the values of the other qubits.

Based on the query complexity of the algorithm in Rattew & Reberst (2023) to the state preparation unitary,  $O(\log(1/\varepsilon))$  rounds will suffice to obtain a state

$$|\tilde{\varphi}_\sigma\rangle = \alpha |0\rangle^{\otimes n+4} |\tilde{y}\rangle + |\phi\rangle, \quad (\text{B.26})$$

such that

$$\left\| |\tilde{y}\rangle - \left| \frac{1}{\|\sigma(U_1 x)\|_2} \sigma(U_1 x) \right\rangle \right\|_2 < \varepsilon. \quad (\text{B.27})$$

Bob then applies  $U_2$  to the state  $|\tilde{\varphi}_\sigma\rangle$  conditioned on the first  $n + 4$  qubits being in the state  $|0\rangle^{\otimes n+4}$ . The state  $|\phi\rangle$  is unaffected. Unitary of  $U_2$  combined with the above bound guarantees

$$\left\| |\hat{y}\rangle - \left| U_2 \frac{1}{\|\sigma(U_1 x)\|_2} \sigma(U_1 x) \right\rangle \right\|_2 < \varepsilon. \quad (\text{B.28})$$

Additionally, from Theorem 3 of Rattew & Rebertost (2023) we are guaranteed that  $\alpha = O(1)$ .  $\square$

## C DATA PARALLELISM

Data parallelism involves storing multiple copies of a model on different devices and training each copy on a subset of the full data. We consider a model of the form

$$|\varphi(\Theta, x)\rangle \equiv \left( \prod_{\ell=L}^1 U_\ell(\theta_\ell, x) \right) |x\rangle, \quad (\text{C.1})$$

where  $x$  is an  $N_1 \times N_2$  matrix which we write as  $x = [x_A, x_B]$  for two  $N_1/2 \times N_2$  matrices  $x_A, x_B$ . Assume also that  $\|x\|_F = 1$ . This model can be used to define a distributed problem with data parallelism by considering the following inputs to both players:

$$\begin{aligned} \text{Alice : } & x_A, \{U_\ell\}, \\ \text{Bob : } & x_B, \{U_\ell\}. \end{aligned} \quad (\text{C.2})$$

The state  $|x\rangle$  can be prepared in a single round of communication involving  $\log(N_1 N_2)$  qubits. Alice simply prepares the state

$$\begin{aligned} |x_A\rangle + \sqrt{1 - \|x_A\|_F^2} |N_1/2, 0\rangle &= (x_A)_{ij} \sum_{i=0}^{N_1/2-1} \sum_{j=0}^{N_2-1} |i, j\rangle + \sqrt{1 - \|x_A\|_F^2} |N_1/2, 0\rangle \\ &= (x_A)_{ij} \sum_{i=0}^{N_1/2-1} \sum_{j=0}^{N_2-1} |i, j\rangle + \|x_B\|_F |N_1/2, 0\rangle, \end{aligned} \quad (\text{C.3})$$

using zero-based indexing of the elements of  $x_A$ . After sending this to Bob, he applies the unitary

$$\frac{1}{\|x_B\|_F} (x_B)_{i,j} \sum_{i=0}^{N_1/2-1} \sum_{j=0}^{N_2-1} |ij\rangle \langle N_1/2, 0| + h.c.. \quad (\text{C.4})$$

The resulting state is  $|x\rangle$ . As before, the gradients with respect to the parameters of the unitaries  $\{U_\ell\}$  can be estimated by preparing copies of this state and using shadow tomography. The number of copies will again be logarithmic in  $N_1, N_2$  and the number of trainable parameters.

## D ADVANTAGES IN END-TO-END TRAINING

In the following, we make use of well-known convergence rates for stochastic gradient descent:

**Lemma 6** (Bubeck (2014)). *Given an objective function  $\mathcal{L}(\Theta)$  with a minimum at  $\Theta^*$  and a stochastic gradient oracle that returns a noisy estimate of the gradient  $g(\Theta)$  such that  $\mathbb{E}g(\Theta) = \nabla\mathcal{L}(\Theta)$ ,  $\mathbb{E}\|g\|_2^2 \leq G^2$ , and denoting by  $\Theta^{(0)}$  a point in parameter space and  $R = \|\Theta^{(0)} - \Theta^*\|_2$ , we have:*

i) *If  $\mathcal{L}$  is convex in a Euclidean ball of radius  $R$  around  $\Theta^*$ , then gradient descent with step size  $\eta = \frac{R}{G}\sqrt{\frac{2}{T}}$  achieves*

$$\mathbb{E}\mathcal{L}\left(\frac{1}{T}\sum_{t=1}^T \Theta^{(t)}\right) - \mathcal{L}(\Theta^*) \leq RG\sqrt{\frac{2}{T}}. \quad (\text{D.1})$$

ii) *If  $\mathcal{L}$  is  $\lambda$ -strongly convex in a Euclidean ball of radius  $R$  around  $\Theta^*$ , then gradient descent with step size  $\eta_t = \frac{2}{\lambda(t+1)}$  achieves*

$$\mathbb{E}\mathcal{L}\left(\frac{1}{T(T+1)}\sum_{t=1}^T 2t\Theta^{(t)}\right) - \mathcal{L}(\Theta^*) \leq \frac{2G^2}{\lambda(T+1)}. \quad (\text{D.2})$$

## D.1 DISTRIBUTED PROBABILISTIC COORDINATE DESCENT

---

### Algorithm 1 Distributed Probabilistic Coordinate Descent

---

**Input:** Alice:  $x, \{A_\ell\}, \Theta_A, \{\eta_t\}, T$ . Bob:  $\{B_\ell\}, \Theta_B, \{\eta_t\}, T$ .

**Output:** Alice: Updated parameters  $\Theta_A^{(T)}$ . Bob: Updated parameters  $\Theta_B^{(T)}$ .

- 1: Alice and Bob each pre-process their coefficient vectors  $\beta_A, \beta_B$  to enable efficient sampling.
  - 2: Alice sends  $\|\beta_A\|_1$  to Bob.  $\triangleright O(\log P)$  bits of classical communication.
  - 3: **for**  $t \in \{1, \dots, T\}$  **do**
  - 4:     Bob samples  $b \sim \text{Bernoulli}(\|\beta_A\|_1 / \|\beta\|_1)$  and sends  $b$  to Alice  $\triangleright 1$  bit of classical communication.
  - 5:     **if**  $b == 0$  **then**
  - 6:         Bob samples  $(\ell, i)$  from the discrete distribution defined by  $\text{abs}(\beta_B)$
  - 7:         Bob create the state  $|\psi_{\ell 0}^B\rangle$   $\triangleright O(L)$  rounds of quantum communication
  - 8:         Bob measures  $\hat{E}_{\ell i}^B$ , as defined in eq. (D.3), obtaining a result  $m \in \{-1, 1\}$
  - 9:         Bob sets  $\theta_{\ell i}^B \leftarrow \theta_{\ell i}^B - \eta_t \text{sign}(\beta_{\ell i}^B) \|\beta\|_1 m$
  - 10:     **else**
  - 11:         Alice runs steps 6-9, (replacing  $B$  with  $A$ )
  - 12:     **end if**
  - 13: **end for**
- 

Given distributed states of the form eq. (4.1), optimization over  $\Theta$  can be performed using Algorithm 1. We verify the correctness of this algorithm and provide convergence rates following Harrow & Napp (2021). Define the Hermitian measurement operator

$$\hat{E}_{\ell i}^Q = \left( |0\rangle\langle 0| \otimes I - i |1\rangle\langle 1| \otimes \mathcal{P}_{\ell i}^Q \right)^\dagger X_a \left( |0\rangle\langle 0| \otimes I - i |1\rangle\langle 1| \otimes \mathcal{P}_{\ell i}^Q \right), \quad (\text{D.3})$$

with eigenvalues in  $\{-1, 1\}$ . Note that  $\beta_{\ell i}^Q \langle \psi_{\ell 0}^Q | \hat{E}_{\ell i}^Q | \psi_{\ell 0}^Q \rangle = \frac{\partial \mathcal{L}}{\partial \theta_{\ell i}^Q}$ , and this is essentially a compact way of representing a Hadamard test for the relevant expectation value. Now consider a gradient estimator that first samples  $(Q, \ell, i)$  with probability  $|\beta_{\ell i}^Q| / \|\beta\|_1$ , then returns a one-sparse vector with  $g_{\ell i}^Q = \text{sign}(\beta_{\ell i}^Q) \|\beta\|_1 m$ , where  $m$  is the result of a single measurement of  $\hat{E}_{\ell i}^Q$  using the state  $|\psi_{\ell 0}^Q\rangle$ . For this estimator we have

$$\mathbb{E} g_{\ell i}^Q = \text{sign}(\beta_{\ell i}^Q) \|\beta\|_1 \frac{|\beta_{\ell i}^Q|}{\|\beta\|_1} \langle \psi_{\ell 0}^A | \hat{E}_{\ell i}^Q | \psi_{\ell 0}^A \rangle = \frac{\partial \mathcal{L}}{\partial \theta_{\ell i}^Q}, \quad (\text{D.4})$$

where the expectation is taken over both the index sampling process and the quantum measurement. The procedure generates a valid gradient estimator.

In order to show convergence, one simply notes that by construction,  $\|g\|_2 = \|\beta\|_1$ . It then follows immediately from Lemma 6 that, with an appropriately chosen step size, Algorithm 1 achieves  $\mathbb{E} \mathcal{L}(\Theta) - \mathcal{L}(\Theta^*) \leq \varepsilon_0$  for a convex  $\mathcal{L}$  using

$$\frac{2 \|\Theta^{(0)} - \Theta^*\|_2^2 \|\beta\|_1^2}{\varepsilon_0^2} \quad (\text{D.5})$$

queries. For a  $\lambda$ -strongly convex  $\mathcal{L}$ , only

$$\frac{2 \|\beta\|_1^2}{\lambda \varepsilon_0} + 1 \quad (\text{D.6})$$

queries are required. The pre-processing in step 1 of Algorithm 1 requires time  $O(P \log P)$  and subsequently enables sampling in time  $O(1)$  using e.g. Walker (1974)<sup>4</sup>.

---

**Algorithm 2** Shadow Tomographic Distributed Gradient Descent

---

**Input:** Alice:  $x, \{A_\ell\}, \Theta_A, \eta, T$ . Bob:  $\{B_\ell\}, \Theta_B, \eta, T$ .

**Output:** Alice: Updated parameters  $\Theta_A^{(T)}$ . Bob: Updated parameters  $\Theta_B^{(T)}$ .

```
1: for  $t \in \{1, \dots, T\}$  do
2:   for  $\ell \in \{1, \dots, L\}$  do
3:     Alice creates  $\tilde{O}(\log^4 P \log D \log(L/\delta)/\varepsilon^4)$  copies of  $|\psi_{\ell 0}^A(\Theta^{(t)})\rangle$   $\triangleright O(L)$  rounds of communication
4:     Alice runs Shadow Tomography to estimate  $\{\mathbb{E}E_{\ell i}^A(\Theta^{(t)})\}$  up to error  $\varepsilon$ , denoting these  $g_{\ell i}^A(\Theta^{(t)})$ .
5:     Bob creates  $\tilde{O}(\log^4 P \log D \log(L/\delta)/\varepsilon^4)$  copies of  $|\psi_{\ell 0}^B(\Theta^{(t)})\rangle$   $\triangleright O(L)$  rounds of communication
6:     Bob runs Shadow Tomography to estimate  $\{\mathbb{E}E_{\ell i}^B(\Theta^{(t)})\}$  up to error  $\varepsilon$ , denoting these  $g_{\ell i}^B(\Theta^{(t)})$ .
7:     Alice sets  $\theta_{\ell i}^{A(t+1)} \leftarrow \theta_{\ell i}^{A(t)} - \eta g_{\ell i}^A(\Theta^{(t)})$ .
8:     Bob sets  $\theta_{\ell i}^{B(t+1)} \leftarrow \theta_{\ell i}^{B(t)} - \eta g_{\ell i}^B(\Theta^{(t)})$ .
9:   end for
10: end for
```

---

## D.2 SHADOW TOMOGRAPHIC DISTRIBUTED GRADIENT DESCENT

## E COMMUNICATION COMPLEXITY OF LINEAR CLASSIFICATION

We consider the following classification problem:

**Problem 5** (Distributed Linear Classification). *Alice and Bob are given  $x, y \in S^N$ , with the promise that  $|x \cdot y| \geq \gamma$  for some  $0 \leq \gamma \leq 1$ . Their goal is to determine the sign of  $x \cdot y$ .*

This is one of the simplest distributed inference problem in high dimensions that one can formulate.  $x$  can be thought of as the input to the model, while  $y$  defines a separating hyperplane with some margin. Since with finite margin we are only required to resolve the inner product between the vectors to some finite precision, it might seem that an exponential quantum advantage should be possible for this problem by encoding the inputs in the amplitudes of a quantum state. However, we show that classical algorithms can leverage this structure as well, and consequently that the quantum advantage in communication that can be achieved for this problem is at most polynomial in  $N$ . We prove this with respect to the randomized classical communication model, in which Alice and Bob are allowed to share random bits that are independent of their inputs<sup>5</sup>.

**Lemma 7.** *The quantum communication complexity of Problem 5 is  $\Omega(\min(\sqrt{1/N}, \sqrt{N/\lceil \gamma N \rceil}))$ . The randomized classical communication complexity of Problem 5 is  $O(\max(N, 1/\gamma^2))$ .*

*Proof.* We first describe a protocol that allows Alice and Bob to solve the linear classification problem with margin  $\gamma$  using  $O(1/\gamma^2)$  bits of classical communication and shared randomness, assuming  $\gamma > 0$ . Note that this bound accords with the notion that the margin rather than the ambient dimension sets the complexity of these types of problems, which is also manifest in the sample complexity of learning with linearly separable data.

Alice and Bob share  $kN$  bits sampled i.i.d. from a uniform distribution over  $\{0, 1\}$ , and that these bits are arranged in a  $k \times N$  matrix  $R$ . Alice and Bob then receive  $x$  and  $y$  respectively, which are valid inputs to the linear classification problem with margin  $\gamma$ . For any  $N$ -dimensional vector  $z$ , define the random projection

$$f : \mathbb{R}^n \rightarrow \mathbb{R}^k, \quad f(z) = \frac{1}{k}(2R - 1)z, \quad (\text{E.1})$$

where addition is element-wise. Applying the Johnson-Lindenstrauss lemma for projections with binary variables Achlioptas (2003), we obtain that if  $k = C/\varepsilon^2$ , for some absolute constant  $C$ , then with probability larger than  $2/3$

---

<sup>4</sup>An even simpler algorithm that sorts the lists as a pre-processing step and uses inverse CDF sampling will enable sampling with cost  $O(\log P)$

<sup>5</sup>This resource can have a dramatic effect on the communication complexity of a problem. The canonical example is equality of  $N$  bit strings, which can be solved with constant success probability using 1 bit of communication and shared randomness, while requiring  $N$  bits of communication otherwise.



we have for any  $z, z' \in \{x, y, 0\}$  (all of these being vectors in  $\mathbb{R}^N$ ),  $f$  is an approximate isometry in the sense

$$(1 - \varepsilon) \|z - z'\|_2^2 \leq \|f(z) - f(z')\|_2^2 \leq (1 + \varepsilon) \|z - z'\|_2^2. \quad (\text{E.2})$$

The key feature of this result is that  $k$  is completely independent of  $N$ . Applying it repeatedly gives

$$\begin{aligned} \|f(x) - f(y)\|_2^2 - \|f(x)\|_2^2 - \|f(y)\|_2^2 &\leq (1 + \varepsilon) \|x - y\|_2^2 - 2(1 - \varepsilon) \\ f(x) \cdot f(y) &\geq (1 + \varepsilon)x \cdot y - 2\varepsilon. \end{aligned} \quad (\text{E.3})$$

Obtaining an upper bound in a similar fashion using the converse inequalities, we have

$$(1 + \varepsilon)x \cdot y - 2\varepsilon \leq f(x) \cdot f(y) \leq (1 - \varepsilon)x \cdot y + 2\varepsilon. \quad (\text{E.4})$$

Assume now that  $x, y$  are valid inputs to the linear classification problem with margin  $\gamma$ , and specifically that  $x \cdot y \geq \gamma$ . The lower bound above gives

$$(1 + \varepsilon)\gamma - 2\varepsilon \leq f(x) \cdot f(y), \quad (\text{E.5})$$

and if we choose  $\varepsilon = \gamma/8$  we obtain

$$\frac{\gamma}{2} \leq (1 + \frac{\gamma}{8})\gamma - \frac{\gamma}{4} \leq f(x) \cdot f(y), \quad (\text{E.6})$$

where we used  $\gamma \leq 1$ . Similarly, if instead  $x \cdot y \leq -\gamma$  we obtain

$$f(x) \cdot f(y) \leq -(1 - \frac{\gamma}{8})\gamma + \frac{\gamma}{4} \leq -\frac{\gamma}{2}. \quad (\text{E.7})$$

It follows that if Alice computes  $f(x)$  and sends the resulting  $O(k) = O(1/\gamma^2)$  bits that describe this vector to Bob (assuming some finite precision that is large enough so as not to affect the margin, which will contribute ), Bob can simply compute  $f(x) \cdot f(y)$  which will reveal the result of the classification problem, which he can then communicate to Alice using a single bit.

If  $\gamma = 0$  there is a trivial  $O(N)$  classical algorithm where Alice sends Bob  $x$ .

We next describe the quantum lower bound for Problem 5. Denote by  $d_H$  the Hamming distance between binary vectors. We will use lower bounds for the following problem:

**Problem 6** (Gap Hamming with general gap). *Alice and Bob are given  $\hat{x}, \hat{y} \in \{0, 1\}^N$  respectively. Given a promise that either  $d_H(\hat{x}, \hat{y}) \geq N/2 + g/2$  or  $d_H(\hat{x}, \hat{y}) \leq N/2 - g/2$ , Alice and Bob must determine which one is the case.*

There is a simple reduction from Problem 6 to Problem 5 for certain values of  $\gamma$ , which we will then use to obtain a result for all  $\gamma$ . Assuming Alice is given  $\hat{x}$  and Bob is given  $\hat{y}$ , they construct unit norm real vectors by  $x = (2\hat{x} - 1)/\sqrt{N}, y = (2\hat{y} - 1)/\sqrt{N}$  with addition performed element-wise.

If  $d_H(x, y) \geq N/2 + g/2$  then

$$\begin{aligned} x \cdot y &= \sum_{i, x_i=y_i} \frac{1}{N} + \sum_{i, x_i \neq y_i} \left(-\frac{1}{N}\right) \\ &\geq \frac{N+g}{2} \frac{1}{N} + \frac{N-g}{2} \left(-\frac{1}{N}\right) \\ &= \frac{g}{N}. \end{aligned} \quad (\text{E.8})$$

Similarly,  $d_H(\hat{x}, \hat{y}) \leq N/2 - g/2 \Rightarrow x \cdot y \leq -g/N$ . It follows that  $x, y$  are valid inputs for a linear classification problem over the unit sphere with margin  $2g/N$ . From the results of Nayak & Wu (1998), any quantum algorithm for the Gap Hamming problem with gap  $g \in \{1, \dots, N\}$  requires  $\Omega(\sqrt{N/g})$  qubits of communication. It follows that the linear classification problem requires  $\Omega(\sqrt{1/\gamma})$  qubits of communication. This bound holds for integer  $\gamma N$ . To get a result for general  $0 < \gamma \leq 1$  we simply note that the communication complexity must be a non-decreasing function of  $1/\gamma$ , since any inputs which constitute a valid instance with some  $\gamma$  are also a valid instance for any  $\gamma' < \gamma$ . Given some real  $\gamma$ , the resulting communication problem is at least as hard as the one with margin  $\lceil \gamma N \rceil / N \geq \gamma$ . It follows that a  $\Omega(\sqrt{N/\lceil \gamma N \rceil})$  bound holds for all  $0 < \gamma \leq 1$ .

If  $\gamma = 0$ , by a similar argument we can apply the lower bound for  $\gamma = 1/N$ , implying that  $\Omega(\sqrt{N})$  qubits of communication are necessary. Once again there is only a polynomial advantage at best.  $\square$

## F EXPRESSIVITY OF QUANTUM CIRCUITS

### F.1 ADDITIONAL RESULTS ON OSCILLATORY FEATURES

Extending the unitaries considered in Section 5 to more than one variable, for two scalar variables  $x, y$  define

$$A_\ell(x) = \text{diag}((e^{-2\pi i \lambda_{\ell 1} x}, \dots, e^{-2\pi i \lambda_{\ell N} x})), \quad (\text{F.1a})$$

$$A_\ell(x, y) = \text{diag}((e^{-2\pi i \lambda_{\ell 1} x}, \dots, e^{-2\pi i \lambda_{\ell N/2} x}, e^{-2\pi i \lambda_{\ell, N/2+1} y}, \dots, e^{-2\pi i \lambda_{\ell N} y})) \quad (\text{F.1b})$$

for  $\ell \in \{1, \dots, L\}$ . Once again  $\{B_\ell\}$  are data-independent unitaries, and we denote by  $|\varphi(x)\rangle, |\varphi(x, y)\rangle$  the states defined by interleaving these unitaries in the manner of eq. (3.1), and by  $\mathcal{L}_1, \mathcal{L}_2$  the corresponding loss functions when measuring  $X_0$ .

While the circuits in Section 5 enable one to represent a small number of frequencies from a set that is exponential in  $L$ , one can easily construct circuits that are supported on an exponentially large number of frequencies, as detailed in Lemma 8. We also use measures of expressivity of classical neural networks known as *separation rank* to show that circuits within the class eq. (3.1) can represent complex correlations between their inputs. For a function  $f$  of two variables  $y, z$ , its separation rank is defined by

$$\text{sep}(f) \equiv \min \left\{ R : f(x) = \sum_{i=1}^R g_i(y) h_i(z) \right\}. \quad (\text{F.2})$$

If for example  $f$  cannot represent any correlations between  $y$  and  $z$ , then  $\text{sep}(f) = 1$ . When computed for certain classes of neural networks,  $y, z$  are taken to be subsets of a high-dimensional input. The separation rank can be used for example to quantify the inductive bias of convolutional networks towards learning local correlations Cohen & Shashua (2016), the effect of depth in recurrent networks Levine et al. (2017), and the ability of transformers to capture correlations across sequences as a function of their depth and width Levine et al. (2020).

We find that the output of estimating an observable using circuits of the form eq. (F.1) can be supported on an exponential number of frequencies, and consequently has a large separation rank:

**Lemma 8.** *For  $\{\lambda_{\ell i}\}$  drawn i.i.d. from any continuous distribution and parameterized unitaries  $\{B_\ell\}$  such that the real and imaginary parts of each entry in these matrices is a real analytic function of parameters  $\Theta$  drawn from a subset of  $\mathbb{R}^{P^L}$ , aside from a set of measure 0 over the choice of  $\{\lambda_{\ell i}\}, \{B_\ell\}$ ,*

i) *The number of nonzero Fourier components in  $\mathcal{L}_1$  is  $\left(\frac{N(N-1)}{2}\right)^{L-1} N$ .*

ii)

$$\text{sep}(\mathcal{L}_2) = 2 \left(\frac{N(N-1)}{2}\right)^{L-1} N. \quad (\text{F.3})$$

Proof: Appendix B

This almost saturates the upper bound on the number of frequencies that can be expressed by a circuit of this form that is given in Schuld et al. (2020). The separation rank implies that complex correlations between different parts of the sequence can in principle be represented by such circuits. The constraint on  $\{B_\ell\}$  is quite mild, and applies to standard choices of parameterize unitaries.

The main shortcoming of a result such as Lemma 8 is that it is not robust to measurement error as it is based on constructing states that are equal weight superpositions of an exponential number of terms. It is straightforward to show that circuits of this form can serve as universal function approximators, at least for a small number of variables. For high-dimensional functions it is unclear when a communication advantage is possible, as we describe below.

**Lemma 9.** *Let  $f$  be a  $p$ -times continuously differentiable function with period 1, and denote by  $\hat{f}_{:M}$  the vector of the first  $M$  Fourier components of  $f$ . If  $\|\hat{f}_{:M}\|_1 = 1$  then there exists a circuit of the form eq. (3.1) over  $O(\log M)$  qubits such that*

$$\|\mathcal{L} - f\|_\infty \leq \frac{C}{M^{p-1/2}} \quad (\text{F.4})$$

for some absolute constant  $C$ .

Proof: Appendix B

This result improves upon the result in Pérez-Salinas et al. (2019); Schuld et al. (2020) about universal approximation with similarly structured circuits both because it is non-asymptotic and because it shows uniform convergence rather than convergence in  $L_2$ . Non-asymptotic results universal approximation results were also obtained recently by Gonon & Jacquier (2023), however their approximation error scales polynomially with the number of qubits, as opposed to exponentially as in Lemma 9.

The result of Lemma 9 applies to an  $L = 1$  circuit. The special hierarchical structure of the Fourier transform implies that the same result can be obtained using even simpler circuits with larger  $L$ . Consider instead single-qubit data-dependent unitaries over  $L + 1$  qubits that take the form

$$A_\ell = |0\rangle_0 \langle 0|_0 + |1\rangle_0 \langle 1|_0 \left( |0\rangle_{\ell+1} \langle 0|_{\ell+1} + e^{2\pi i 2^{\ell-1} x} |1\rangle_{\ell+1} \langle 1|_{\ell+1} \right), \quad (\text{F.5})$$

for  $\ell \in \{1, \dots, L\}$ . This is simply a single term in a hierarchical decomposition of the same feature matrix we had in the shallow case, since

$$\prod_{\ell=1}^L A_\ell = |0\rangle_0 \langle 0|_0 \otimes I_{1:L} + |1\rangle_0 \langle 1|_0 \otimes I_1 \otimes \left( \sum_{m=0}^{2^L-1} e^{2\pi i m x} |m\rangle \langle m| \right), \quad (\text{F.6})$$

which is identical to eq. (B.20). As before, set

$$B_1 = |\hat{f}\rangle \langle 0| + |0\rangle \langle \hat{f}|, \quad (\text{F.7})$$

with  $N/4 = 2^L$  and  $B_\ell = I$  for  $\ell > 1$ . This again gives an approximation of  $f$  up to normalization. The data-dependent unitaries are particularly simple when decomposed in this way. The fact that they act on a single qubit and thus have "small width" is reminiscent of classical depth-separation result such as Cohen et al. (2015), where it is shown that (roughly speaking) within certain classes of neural network, in order to represent the function implemented by a network of depth  $L$ , a shallow network must have width exponential in  $L$ . In this setting as well the expressive power as measured by the convergence rate of the approximation error grows exponentially with  $L$  by eq. (B.19).

The circuits above can be generalized in a straightforward way to multivariate functions of the form  $f : [-1/2, 1/2]^D \rightarrow \mathbb{R}$  and combined with multivariate generalization of eq. (B.19). In this case the scalar  $m$  is replaced by a  $D$ -dimensional vector taking  $M^D$  possible values, and we can define

$$A_1(x) = |0\rangle_0 \langle 0|_0 \otimes I_{1:D \log M-1} + |1\rangle_0 \langle 1|_0 \otimes I_1 \otimes \left( \sum_{m \in [M]^D} e^{2\pi i m \cdot x} |m\rangle \langle m| \right). \quad (\text{F.8})$$

Note that using this feature map, the number of neurons is linear in the spatial dimension  $D$ . Because of this, such circuits are not strictly of the form eq. (3.1) since it is not the case that  $M' = O(N)$  where  $N'$  is the Hilbert space on which the unitaries in the circuit act and  $N$  is the input size. An alternative setting where the features themselves are also learned from data could enable much more efficient approximation of functions that are sparse in Fourier space.

## F.2 UNITARY TRANSFORMERS

Transformers based on self-attention Vaswani et al. (2017) form the backbone of large language models Brown et al. (2020); Barham et al. (2022) and foundation models more generally Bommasani et al. (2021). A self-attention layer, which is the central component of transformers, is a map between sequences in  $\mathbb{R}^{S \times N}$  (where  $S$  is the sequence length) defined in terms of weight matrices  $W_Q, W_K, W_V \in \mathbb{R}^{N \times N}$ , given by

$$X'(X) = \text{softmax} \left( \frac{XW_Q W_K^T X^T}{\sqrt{N}} \right) XW_V \equiv A(X)XW_V, \quad (\text{F.9})$$

where  $\text{softmax}(x)_i = e^{x_i} / \sum_j e^{x_j}$  for a vector  $x$ , and acts row-wise on matrices. There is an extensive literature on replacing the softmax-based attention matrix  $A(X)$  with matrices that can be computed more efficiently, which can markedly improve the time complexity of inference and training without a significant effect on performance Katharopoulos et al. (2020); Levine et al. (2020). In some cases  $A(X)$  is replaced by a unitary matrix Lee-Thorp et al. (2021). Remarkably, recent work shows that models without softmax layers can in fact outperform standard transformers on benchmark tasks while enabling faster inference and a reduced memory footprint Sun et al. (2023).

Considering a simplified model that does not contain the softmax operation as in Levine et al. (2020) and dropping normalization factors, the linear attention map is given by

$$X'_{\text{lin}}(X) = XW_QW_K^T X^T XW_V. \quad (\text{F.10})$$

Iterating this map twice gives

$$X'_{\text{lin}}(X'_{\text{lin}}(X)) = XW_Q^{(1)}W_K^{(1)T} X^T XW_V^{(1)}W_Q^{(2)}W_K^{(2)T}W_V^{(1)T} X^T XW_K^{(1)}W_Q^{(1)T} X^T XW_Q^{(1)}W_K^{(1)T} X^T XW_V^{(1)}W_V^{(2)}. \quad (\text{F.11})$$

Iterating this map  $K$  times (with different weight matrices at each layer) gives a function of the form:

$$X_{\text{lin}}^{(K)}(X) = XR_0 \prod_{\ell=1}^{(3^K-1)/2} (X^T X R_\ell), \quad (\text{F.12})$$

where the  $\{R_\ell\}$  matrices depend only on the trainable parameters. If we now constrain these to be parameterized unitary matrices, and additionally replace  $X^T X$  with a unitary matrix  $U_X$  encoding features of the input sequence itself, then the  $i$ -th row of the output of this model is encoded in the amplitudes of a state of the form eq. (3.6) with  $L = (3^K - 1)/2 + 1$ ,  $|\psi(x)\rangle = |X_i\rangle$ ,  $A_\ell(x) = U_X$ ,  $B_\ell = R_\ell$ .

### F.3 ENSEMBLING AND POINT-WISE NONLINEARITIES

An additional method for increasing expressivity while maintaining an advantage in communication is through ensembling. Given  $K$  models of the form Definition 3.1 with  $P$  parameters each, one can combine their loss functions  $\mathcal{L}_1, \dots, \mathcal{L}_K$  into any differentiable nonlinear function

$$\tilde{\mathcal{L}}(\mathcal{L}_1(\Theta_1, x), \dots, \mathcal{L}_K(\Theta_K, x), \tilde{\Theta}, x) \quad (\text{F.13})$$

that depends on additional parameters  $\tilde{\Theta}$ . As long as  $K$  and  $|\tilde{\Theta}|$  scale subpolynomially with  $N$  and  $P$ , the gradients for this more expressive model can be computed while maintaining the exponential communication advantage in terms of  $N, P$ .

## G REALIZING QUANTUM COMMUNICATION

Given the formidable engineering challenges in building a large, fault tolerant quantum processor Arute et al. (2019); Google Quantum AI (2023), the problem of exchanging coherent quantum states between such processors might seem even more ambitious. We briefly outline the main problems that need to be solved in order to realize quantum communication and the state of progress in this area, suggesting that this may not be the case.

We first note that sending a quantum state between two processors can be achieved by the well-known protocol of quantum state teleportation Bennett et al. (1993); Gordon & Rigolin (2005). Given an  $n$  qubit state  $|\psi\rangle$ , Alice can send  $|\psi\rangle$  to Bob by first sharing  $n$  Bell pairs of the form

$$|b\rangle = \frac{1}{\sqrt{2}} (|0\rangle|0\rangle + |1\rangle|1\rangle), \quad (\text{G.1})$$

(sharing such a state involves sending a one of the two qubits to Bob) and subsequently performing local processing on the Bell pairs and exchanging  $n$  bits of classical communication. Thus quantum communication can be reduced to communicating Bell pairs up to a logarithmic overhead, and does not require say transmitting an arbitrary quantum state in a fault tolerant manner, which appears to be a daunting challenge given the difficulty of realizing quantum memory on a single processor. Bell pairs can be distributed by a third party using one-way communication.

In order to perform quantum teleportation, the Bell pairs must have high fidelity. As long as the fidelity of the communicated Bell pairs is above .5, purification can be used produce high fidelity Bell pairs Bennett et al. (1995), with the fidelity of the purified Bell pair increasing exponentially with the number of pairs used. Thus communicating arbitrary quantum states can be reduced to communicating noisy Bell pairs.

Bell pair distribution has been demonstrated across multiple hardware platforms including superconducting waveguides Magnard et al. (2020), optical fibers Krutyanskiy et al. (2022), free space optics at distances of over 1,200 kilometers Li et al. (2022). At least in theory, even greater distances can be covered by using quantum repeaters, which span the distance between two network nodes. Distributing a Bell pair between the nodes can then be reduced to sharing Bell pairs only between adjacent repeaters and local processing Azuma et al. (2022).

A major challenge in implementing a quantum network is converting entangled states realized in terms of photons used for communication to states used for computation and vice versa, known as transduction Lauk et al. (2020). Transduction is a difficult problem due to the several orders of magnitude in energy that can separate optical photons from the energy scale of the platform used for computation. Proof of principle experiments have been performed across a number of platforms including trapped ions Krutyanskiy et al. (2022), solid-state systems Pompili et al. (2021), and superconducting qubits operating at microwave frequencies Balram & Srinivasan (2021); Wang et al. (2022).

## H SOME OPEN QUESTIONS

### H.1 EXPRESSIVITY

Circuits that interleave parameterized unitaries with unitaries that encode features of input data are also used in Quantum Signal Processing Low & Chuang (2017); Martyn et al. (2021), where the data-dependent unitaries are time evolution operators with respect to some Hamiltonian of interest. The original QSP algorithm involved a single parameterized rotation at each layer, and it is also known that extending the parameter space from  $U(1)$  to  $SU(2)$  by including an additional rotation improves the complexity of the algorithm and improves its expressivity Motlagh & Wiebe (2023). In both cases however the expressive power (in terms of the degree of the polynomial of the singular values that can be expressed) grows only linearly with the number of interleaved unitaries. Given the natural connection to the distributed learning problems considered here, it is interesting to understand the expressive power of such circuits with more powerful multi-qubit parameterized unitaries.

We present a method of applying a single nonlinearity to a distributed circuit using the results of Rattew & Reberstrost (2023). Since this algorithm requires a state-preparation unitary as input and produces a state with a nonlinearity applied to the amplitudes, it is natural to ask whether it can be applied recursively to produce a state with the output of a deep network with nonlinearities encoded in its amplitudes. This will require extending the results of Rattew & Reberstrost (2023) to handle noisy state-preparation unitaries, yet the effect of errors on compositions of block encodings Chakraborty et al. (2018); Gilyén et al. (2018), upon which these results are based, is relatively well understood. It is also worth noting that these approaches rely on the approximation of nonlinear functions by polynomials, and so it may also be useful to take inspiration directly from classical neural network polynomial activations, which in some settings are known to outperform other types of nonlinearities Michaeli et al. (2023).

### H.2 OPTIMIZATION

The results of Section 4 rely on sublinear convergence rates for general stochastic optimization of convex functions (Lemma 6). It is known however that using additional structure, stochastic gradients can be used to obtain linear convergence (meaning that the error decays exponentially with the number of iterations). This is achievable when subsampling is the source of stochasticity Le Roux et al. (2012), or with occasional access to noiseless gradients in order to implement a variance reduction scheme Johnson & Zhang (2013); Moritz et al. (2016); Gower et al. (2016), neither of which seem applicable to the setting at hand. It is an interesting open question to ascertain whether there is a way to exploit the structure of quantum circuits to obtain linear convergence rates using novel algorithms. Aside from advantages in time complexity, this could imply an exponential advantage in communication for a more general class of circuits.

Conversely, it is also known that given only black-box access to a noisy gradient oracle, an information-theoretic lower bound of  $\Omega(1/T)$  on the error holds given  $T$  oracle queries, precluding linear convergence without additional structure, even for strongly convex objectives Agarwal et al. (2010). Harrow & Napp (2021) provide a similar lower bound for their algorithm, at least for a restricted class of circuits. Perhaps these results be used to show optimality of algorithms that rely on the standard variational circuit optimization paradigm that involves making quantum measurements at every iteration and using these to update the parameters. This might imply that linear convergence is only possible if the entire optimization process is performed coherently.

In this context, we note that the treatment of gradient estimation at every layer and every iteration as an independent shadow tomography problem is likely highly suboptimal, since no use is made of the correlations across iterations between the states and the observables of interest. While in Section 4.2 this is not the case, that algorithm applies only to fine-tuning of a single layer. Is there a way to re-use information between iterations to reduce the resource requirements of gradient descent using shadow tomography? One approach could be warm-starting the classical

resource states by reusing them between iterations. Improvements along these lines might find applications for other problems as well.

### H.3 STRONGER LOWER BOUNDS

The lower bound in Lemma 3 applies to circuits that contain general unitaries, and thus have depth  $\text{poly}(N)$  when compiled using any reasonable set of quantum gates. One can ask whether the lower bound can be strengthened to apply to more restricted classes of unitaries as well, and in particular log-depth unitaries. The rapid growth of complexity of random circuits composed from local gates with depth suggests that this might be possible Brown & Susskind (2017). This is particularly interesting since Algorithm 1 requires only a single measurement per iteration and may thus be suitable for implementation on near-term devices whose coherence times restrict them to implementing  $\log N$  depth circuits.