



An inductive bias from quantum mechanics: learning order effects with non-commuting measurements

Kaitlin Gili^{1,2} · Guillermo Alonso² · Maria Schuld²

Received: 18 January 2024 / Accepted: 23 September 2024
© The Author(s) 2024

Abstract

There are two major approaches to building good machine learning algorithms: feeding lots of data into large models or picking a model class with an “inductive bias” that suits the structure of the data. When taking the second approach as a starting point to design quantum algorithms for machine learning, it is important to understand how mathematical structures in quantum mechanics can lead to useful inductive biases in quantum models. In this work, we bring a collection of theoretical evidence from the quantum cognition literature to the field of quantum machine learning to investigate how non-commutativity of quantum observables can help to learn data with “order effects,” such as the changes in human answering patterns when swapping the order of questions in a survey. We design a multi-task learning setting in which a generative quantum model consisting of sequential learnable measurements can be adapted to a given task — or question order — by changing the order of observables, and we provide artificial datasets inspired by human psychology to carry out our investigation. Our first experimental simulations show that in some cases the quantum model learns more non-commutativity as the amount of order effect present in the data is increased and that the quantum model can learn to generate better samples for unseen question orders when trained on others — both signs that the model architecture suits the task.

Keywords Quantum mechanics · Learning order effects · Non-commuting measurements · Machine learning · Quantum cognition

1 Introduction

In the quest to understand if and how quantum algorithms for machine learning (ML) will be useful, it is valuable to conduct research from more than one angle (Perdomo-Ortiz et al. 2018; Cerezo et al. 2022; Schuld and Killoran 2022). Typical approaches in the past years have investigated learning settings where the computational complexity of the overall algorithm (Hinsche et al. 2021, 2022; Du et al. 2020; Liu et al. 2021; Schreiber et al. 2022; Cai et al. 2022), generalization scores on benchmarks (Gili et al. 2023; Hibat-Allah et al. 2023; Gili et al. 2022a; Moussa et al. 2023), or trainability guarantees (McClellan et al. 2018; Holmes et al. 2022; Wang et al. 2021; Cerezo et al. 2021; Arrasmith et al. 2021; Rudolph et al. 2023) were the primary focus. While these approaches

have indeed pushed our knowledge boundary forward, an alternative angle is to explore the *inductive bias*, or natural structure that models based on quantum mechanics exhibit (Ragone et al. 2023; Pesah et al. 2021; Bowles et al. 2023), and find applications for quantum machine learning that suit this bias. After interesting model classes are found, questions related to classical simulatability and generalization power can be reassessed from a new vantage point. We add to this work here by constructing and studying a quantum model with a bias that results from non-commutativity in quantum measurement.

More specifically, we introduce a learning setting inspired by the literature on “quantum cognition” (Yearsley and Busemeyer 2016; Busemeyer and Bruza 2012; Busemeyer et al. 2011; Wang et al. 2014; Widdows et al. 2023) which allows us to investigate whether the non-commutative nature of quantum observables encourages certain quantum models to learn data distributions that contain order effects. Order effects can be thought of as sequences of random variables whose joint distribution changes when the sequence is permuted; an example that we will consider here is the change

✉ Kaitlin Gili
kaitlinmgili@gmail.com

¹ University of Oxford, Oxford OX1 2JD, UK

² Xanadu, Toronto, ON M5G 2C8, Canada

of yes/no answer patterns to survey questions when humans are presented with questions in different orders. To control the structure in the data (such as the strength and kind of order effect), we create artificial probability distributions of human answering patterns inspired by the human cognition and psychology literature (Moore 2002; Kaplan et al. 2012; Weinberg et al. 2018).

The quantum model we consider to learn distributions over answering patterns is a generative model in which we interpret the results of a trainable two-outcome measurement as the answer to a question. An answer pattern is then a collection of N yes/no answers generated by measuring N observables in sequence. The model gets adapted to a given question order by permuting the observables accordingly. Mathematically and conceptually, this investigation design differs fundamentally from existing generative models such as quantum Born machines (Benedetti et al. 2019; Liu and Wang 2018b; Cheng et al. 2018; Gili et al. 2022b) and quantum Boltzmann machines (Kieferova and Wiebe 2017; Korenkevych et al. 2016; Cheng et al. 2018). Firstly, instead of measuring the state of a set of qubits, we take subsequent measurements of the same state, which can be understood as “collapsing” the wavefunction several times. Since such a workflow exceeds the typical blueprint of quantum algorithms, one of our contributions is the implementation of the model on a mixed-state simulator without access to sequential measurements, for which we use PennyLane (Bergholm et al. 2022). Secondly, the problem we consider is a multi-task learning problem (Ruder 2017; Bao et al. 2021), where the goal is to generalize to unseen training data *as well as* to related unseen tasks. Bowles et al. (2023) independently found that such a framework suits as a bridge between quantum foundations and machine learning.

We use this study design to walk the first steps of exploring two basic questions:

1. As we increase the strength of the order effect in the dataset, does the model learn a stronger degree of non-commutativity?
2. Can the quantum model generalize better on unseen orders after training on other orders?

An affirmation of both — and we will define the criteria in more detail below — would give support to the hypothesis that non-commutativity can create an inductive bias to learn order effects. Our results presented here indicate that on some of the datasets, the non-commutativity of the trained model grows with the strength of the order effect in the data, and that for up to five observables, training on a sufficiently large number of tasks lowers the generalization error on unseen tasks, even if we never trained on those. We view these results as an initial investigation into whether the inductive bias of non-commutativity is the primary resource driving the model

to learn the underlying pattern of order effects. In our outlook, we conclude with some future research questions prompted from this study that could strengthen or weaken this claim in the future.

We note that the study design does not aim at making any strong statements toward the performance of the quantum model compared to classical machine learning. We also do not test the approach with regard to current industry applications. This may be surprising considering the dominant culture in the quantum machine learning literature, but we believe that investigations with a limited but well-defined scope that study the qualitative behavior of a model provide important additional insights. This will be further discussed in Sect. 4.

2 The quantum learning setting

In the first part of this work, we introduce the learning setting, including the model and problem to solve. We provide a theoretical intuition in Sect. 2 by discussing the nature of non-commutativity in quantum theory and its relation to structures that exist in binary distributions containing order effects — i.e., measured changes in the distribution due to the ordering of the variables. Here, we also define a measure for the *amount of non-commutativity*, which we utilize to quantify the non-commutativity within the model at each training step. Building off of these insights, in Sect. 2.1, we introduce a multi-task quantum generative architecture that can be used to investigate whether the structure of non-commutativity biases the model toward learning binary order effects and provide an overview of the model’s implementation. Lastly, we introduce the psychology-inspired datasets intentionally designed to carry out the study.

Motivating the model design A critical motivation of this investigation is taken from a result that is increasingly recognized in the classical machine learning literature: models that contain an *inductive bias* that matches the structure in the dataset it aims to learn, will ultimately have better generalization performance with typically fewer samples (Goyal and Bengio 2022; Haussler 1988; Feinman and Lake 2018). Informally, an inductive bias can be defined as the preference of a model for one hypothesis that is consistent with the training data over another. Inductive biases encourage the model toward this hypothesis, including extra information about the problem in the model itself, rather than in the training data. Examples of inductive biases in classical machine learning are priors in Bayesian networks, translation-invariance in convolutional neural networks (CNNs), or regularizers in linear regression (Goyal and Bengio 2022; Cohen and Shashua 2017; O’Shea and Nash 2015). Turning to the induc-

tive biases of quantum models, we can leverage non-classical mathematical structures and behaviors that exist in quantum mechanics, and search for alignments within correlations in datasets (Bowles et al. 2023). Here, we discuss non-commuting observable measurements in quantum mechanics and provide theoretical intuition for its structural connection to binary distributions with order effects. This choice is inspired by a sub-field known as “quantum cognition” which utilizes elements from quantum theory to build models that explain human cognitive behavior (Yearsley and Busemeyer 2016; Busemeyer and Bruza 2012; Busemeyer et al. 2011; Wang et al. 2014). The basic hypothesis of quantum cognition is that certain features of distributions generated by human decision-making are difficult to model by classical probability theory while naturally found in the statistics of quantum mechanics (Wang et al. (2014)). An example (Pothos and Busemeyer 2022) is the *conjunction fallacy*, or the consistent observation that humans estimate the probability of two conditions to be true — such as the probability that a Scandinavian person has blond hair *and* blue eyes — higher than the probability of one of them to be true (i.e., the probability that a Scandinavian person has blond hair). This violates the set-theoretic foundations of classical probability theory which devise $p(A \& B) \leq p(A)$. Order effects are another example that has been studied using small-scale datasets of medical and legal decision-making (Trueblood and Busemeyer 2011). The quantum cognition literature concludes that quantum mechanics could be an interesting mathematical framework to model such effects. In some sense, our study can be seen as an attempt to probe this observation from a machine learning angle.

2.1 Non-commutativity

Let us review the basics of sequential measurements in quantum theory in order to motivate the design of our quantum model, a section that readers with a background in quantum theory can easily skip. Consider a set of N quantum observables $\{Q_n\}$, $n = 1, \dots, N$ describing measurements that can have two possible outcomes. Note that there are $N!$ permutations of these measurements which we call “(measurement) orders.” Intuitively, one can view each observable Q_n to be a question regarding binary information that we would like to know about the quantum system. Without loss of generality, we assume that each observable Q_n of arbitrary size only contains the eigenvalues $\{-1, 1\}$ and denote by $\{|u_1^n\rangle, \dots, |u_m^n\rangle\}$ the eigenstates associated with $+1$ and by $\{|v_1^n\rangle, \dots, |v_l^n\rangle\}$ the eigenstates associated with -1 . Let $|\psi_0\rangle$ be a quantum state that we wish to measure with respect to the observable, such that we can express it using the eigenbasis of Q_n :

$$|\psi_0\rangle = \sum_i \alpha_i |u_i^n\rangle + \sum_i \beta_i |v_i^n\rangle \tag{1}$$

After the measurement the state $|\psi_1\rangle$ of the system is given by $\frac{1}{\sqrt{\sum_i |\alpha_i|^2}} \sum_i \alpha_i |u_i^n\rangle$ with probability $\sum_i |\alpha_i|^2$ and by $\frac{1}{\sqrt{\sum_i |\beta_i|^2}} \sum_i \beta_i |v_i^n\rangle$ with probability $\sum_i |\beta_i|^2$.

Now consider a second observable Q_m , and express the eigenstates of Q_n in this new eigenbasis, $|u_i^n\rangle = \sum_j \gamma_{ij}^u |u_j^m\rangle + \sum_j \zeta_{ij}^u |v_j^m\rangle$, and $|v_i^n\rangle = \sum_j \gamma_{ij}^v |u_j^m\rangle + \sum_j \zeta_{ij}^v |v_j^m\rangle$. This allows us to rewrite $|\psi_1\rangle$ in the eigenbasis of Q_m , and — after cancelling some terms — obtain the following probability distribution over the two measurement outcomes:

$$P_{nm}(+1, +1) = \sum_j \sum_i \left| \alpha_i \gamma_{ij}^u \right|^2 \tag{2}$$

$$P_{nm}(-1, +1) = \sum_j \sum_i \left| \beta_i \gamma_{ij}^v \right|^2 \tag{3}$$

$$P_{nm}(+1, -1) = \sum_j \sum_i \left| \alpha_i \zeta_{ij}^u \right|^2 \tag{4}$$

$$P_{nm}(-1, -1) = \sum_j \sum_i \left| \beta_i \zeta_{ij}^v \right|^2. \tag{5}$$

If the two Hermitian observables Q_n, Q_m share common eigenspaces, they can be written as $Q_n = V D_n V^\dagger$ and $Q_m = V D_m V^\dagger$, where V is a unitary and $D_{\{n,m\}}$ a diagonal operator. It then follows that

$$Q_n Q_m = V D_n V^\dagger V D_m V^\dagger = V D_n D_m V^\dagger = V D_m D_n V^\dagger = V D_m V^\dagger V D_n V^\dagger = Q_m Q_n \tag{6}$$

When the observables do not share the same eigenstates, the above terms will not cancel, and the non-commutativity relation $Q_n Q_m - Q_m Q_n \neq 0$ holds. As $Q_n Q_m \neq Q_m Q_n$, the distribution in (2) that results from measuring $Q_n Q_m$ on $|\psi_0\rangle$ will differ from the distribution that results from considering the measurement order $Q_m Q_n$. Hence, the distribution over measurement results depends on the order of the observable measurements. Once again, intuitively, one can consider this as the process of asking a quantum system binary questions in a different order and observing different resulting answer outcomes. Note that for N observable measurements, there exist 2^N combinations of binary answers that the quantum system could reveal and $N!$ order permutations that could produce different binary distributions.

Following (Guo 2016), for N observables in the subset $\Gamma = \{Q_n : 1 \leq n \leq N\}$, the amount of non-commutativity of the total set can be measured by

$$\zeta(\Gamma) := \sum_{n < m} \|[Q_n, Q_m]\|_{\text{Tr}}, \tag{7}$$

where $\|\cdot\|_{\text{Tr}}$ is the trace norm such that $\|Q\|_{\text{Tr}} = \text{Tr}\sqrt{Q^\dagger Q}$. As the models used in this study remain small

in the number of observables, we are able to obtain an exact set of updated observables at each training step and compute this metric directly. Of course, this approach does not scale, and still, it is convenient enough for a small-scale investigation like ours.

Quantum generative model Borrowing intuition from Sect. 2, we put forth an unsupervised quantum generative model that can learn distributions over binary strings $x \in \{-1, 1\}^N$ from measuring a sequence of N quantum observables. We define an unsupervised generative model as one in which given a set of i.i.d. unlabeled datapoints from an unknown target distribution $P(x)$, the goal of the model is to learn the underlying features in the training samples such that the true $P(x)$ can be best approximated, and valuable data can be efficiently generated (Gili et al. 2023; Bengio 2009). Parameterized quantum circuits can be understood as interesting candidates for natural, universal approximators of probability distributions with efficient sampling capabilities, and have been investigated as generators for various ML tasks (Benedetti et al. 2019a, 2018; Rudolph et al. 2022; Gili et al. 2022a, ?; Rudolph et al. 2020).

We provide a visualization of our entire algorithm in Fig. 1. The quantum circuit requires the number of qubits n_{qubits} to grow linearly with the number of observables N in order to have full expressivity over all possible binary

answers. Starting in an initial state $|\psi_0\rangle = |00\dots 0_N\rangle$, we take measurements of subsequent trainable observables, discussed in Sect. 2, each taking the parameterized form $Q_n = V(\theta_n)D_nV(\theta_n)^\dagger$. In the circuit implementation, each unitary transformation $V(\theta_n)$ is composed of arbitrary parameterized quantum gate sequences that play a role in the expressivity of the model. Specifically, we utilize fully expressive single-qubit rotation gates consisting of Pauli- X and Pauli- Z rotations, $R_X(\theta)R_Z(\theta)R_X(\theta)$ respectively on each of the qubits, with $R_l(\theta) = \exp(\frac{-i\theta\sigma_l}{2})$. After each sequence of single-qubit operations, we utilize multi-qubit entangling gates, specifically XX couplers in-between nearest-neighbor qubits, known as a line-topology (Zhu et al. 2019). The gate sequence for each observable Q_n does not change when observables are permuted and only depends on the parameters θ_n . Thus, in our ansatz, the gate count and number of parameterized operations grows linearly with the number of qubits n_{qubits} and the number of observables N . The total number of model parameters in the ansatz is given by $N(4n_{\text{qubits}} - 1)$. Note that understanding the optimal gate sequences for N observables is an open question for future work, especially in the regime that extends beyond binary data.

After implementing $V(\theta_n)$, we take a single (one-shot) measurement of the first qubit in the Pauli- Z basis to obtain the corresponding eigenvalue (binary answer) $x_i \in \{1, -1\}$.

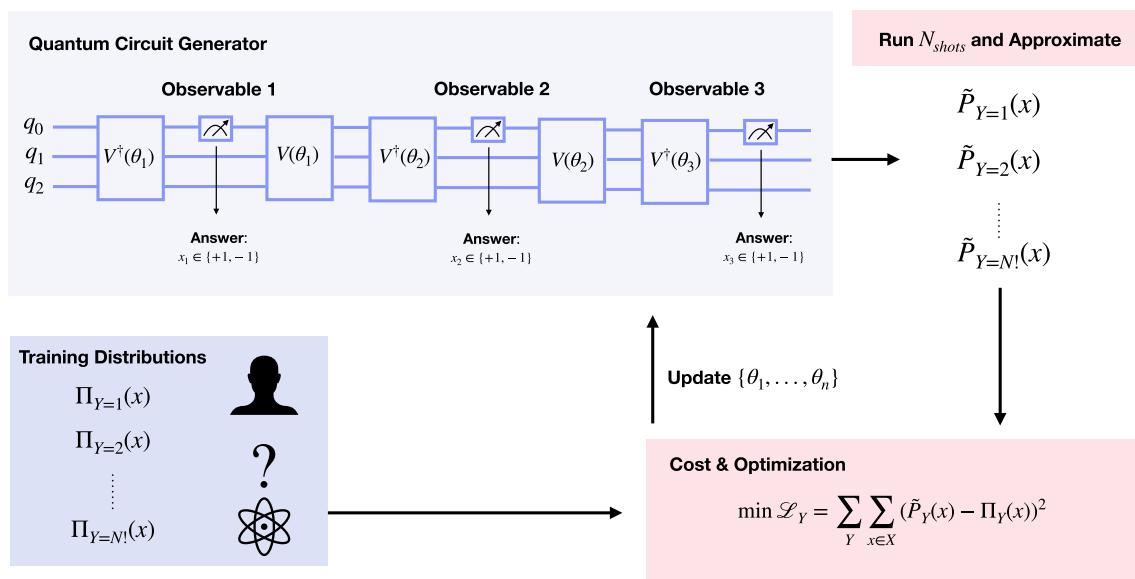


Fig. 1 An example of the multi-task quantum generative architecture for training on binary distributions with three observables. We provide a visual overview of the quantum circuit generator, containing subsequent measurement observables of the form $V(\theta_n)D_nV(\theta_n)^\dagger$. Measurements taken in the Pauli- Z basis allow one to obtain a binary eigenvalue (answer) for each trainable observable. The generator is run for a number of N_{shots} to approximate the desired number of model

output distributions $\tilde{P}_Y(x)$ over binary answer variables, with respect to the order sequence of the observables. Swapping the order sequence simply corresponds to physically swapping the parameters θ_n in the generator, which are indexed by Y . The model is trained to minimize the multi-task loss function \mathcal{L}_Y that computes the LSM loss on the training distribution $\Pi_Y(x)$ and the empirical model distribution $\tilde{P}_Y(x)$ for the desired number of tasks to use for training

While more than one method can be utilized to simulate these sequential measurements, we implement a the circuit using a mixed-state phase damping method that requires an additional qubit in our specific circuit implementation (code provided in 6). The shot samples from our quantum generator are then taken from the distribution $P_Y(x)$, where $Y \in \{1, 2, \dots, N!\}$ refers to the order of the observables (or, equivalently, the order of the ansatz parameters θ_n) in the circuit. Depending on this order, the circuit gives rise to one of the $N!$ probability distributions and can therefore be used to learn different tasks.

The model is trained on a task or order by minimizing the loss function \mathcal{L}_Y between the empirical distribution \tilde{P}_Y generated from the output samples and an empirical distribution over the training data Π_Y . Typical loss functions include distance measures such as the negative log likelihood (NLL) and the maximum mean discrepancy (MMD) loss (Liu and Wang 2018a). In this work, we utilize a simple least mean squares (LMS) loss at each iteration defined as:

$$\mathcal{L}_Y = \sum_{x \in X} (\tilde{P}_Y(x) - \Pi_Y(x))^2, \quad (8)$$

We highlight that the question of how to best train generative quantum models is still open. Computing loss functions that require complete access to the model output probability distribution, like the NLL and LMS above, have exponential scaling relative to the dimension of the data samples and require good parameter initializations. As such, these loss functions are only useful in simple research settings like ours, where the data dimensions remain small. In more practical settings, the current best practice is to utilize loss functions based on f-divergences (Leadbeater et al. 2021) and maximum mean discrepancies (Rudolph et al. 2023).

Once computed, the loss is utilized to approximate the gradient in an Adam optimizer (Zhang 2018), which updates the parameter values and feeds them back into the circuit for the next iteration, until convergence is reached.

Multi-task training scheme Given a set of related learning tasks \mathcal{T} , multi-task learning aims to enhance a model on all tasks by using the knowledge contained in only a few of them Ruder (2017); Bao et al. (2021). The model usually consists of a base model and an adaptation to each task, such as a neural network whose weights are shared between all tasks for all but the final layer. Multi-task learning can be put in the context of generative modeling in the following way: a task $T_i \in \mathcal{T}$ consists of learning to sample from a probability distribution $P_i(x)$ given a finite set of training data from that distribution. The goal is to train on data from several tasks simultaneously with a loss function $\mathcal{L} = \sum_i \mathcal{L}_i(T_i)$ and generate datapoints both from the array of tasks. This multi-task approach is expected to lead to a learning advantage over

training separate architectures — either in terms of run-time, sample complexity, or generalization performance.

In our learning setting, we can define each task T_i as the quantum generator learning a distribution over binary answers with a fixed order of measurement observables. While the observables themselves are used for all tasks, their order — in our case, the parameters of the ansatz — gets adapted to the question order of a specific task. As such, our loss function for training becomes:

$$\mathcal{L}_Y = \sum_Y \sum_{x \in X} (\tilde{P}_Y(x) - \Pi_Y(x))^2 \quad (9)$$

Similarly to training on a single distribution, one can use a gradient descent algorithm to minimize the loss function. Not all orders are required for training, and excluding orders from the training set to use in a test set is necessary to assess generalization performance to new tasks from \mathcal{T} . As the number of tasks required to compute the loss and achieve quality generalization performance for a scaled number of observables remains uncertain, the computational efficiency of the loss function is unclear.

In addition, we want to highlight the distinction here between the model generalizing to *tasks* (out-of-task generalization) vs. the model generalizing to *samples* (out-of-distribution generalization) from each task. In classical multi-task learning, typically very few samples from all tasks are utilized with the goal of obtaining good out-of-distribution generalization (see, for example, (Ishibashi et al. 2021)). Here, we are interested in out-of-task generalization as an evaluation metric to determine whether the bias of non-commutativity in the circuit model helps to learn the general structure of the tasks, i.e., the order effect underlying the distributions. We are using this model and scheme to obtain a more fundamental understanding of how quantum mechanics plays a role in a learning setting and are not trying to claim computational advantage with our quantum model or training scheme.

To summarize our quantum learning setting, we want to emphasize *two main features* that result from the non-commutative nature of our quantum model and training scheme that intuitively argues its effectiveness for learning order effects in data:

1. **A multi-task design:** The model can represent and train on multiple probability distributions at once — specifically those conditioned on the order of observables. Thus, observable order becomes an inherent variable of the model.
2. **A design that matches the data-generating process:** The physical mechanics of the model match the physical process of the data generation. More specifically, obtaining information from a human in a survey corresponds to observing a property of a quantum system.

Psychology-inspired datasets For our artificial datasets, we take inspiration from order effect studies that relate to human cognition and decision-making (Moore 2002; Kaplan et al. 2012; Weinberg et al. 2018). In this context, a classical order effect can be defined as the change in the distribution over human answers based on the order of questions asked. These order effects are prevalent in survey data across most disciplines and within most countries (Stark et al. 2020; Revilla et al. 2016; Weinberg et al. 2018). In the psychology literature, we find many reasons for such order effects to occur. Typically, we see order effects in one of the following circumstances: (1) humans want to appear consistent with their answers such that their responses to previous questions will impact future answers; (2) humans don't want to be redundant so they try to provide new information in sequential answers; (3) humans use direct information or intuition from the previous question in their next answer (known as "priming") (Kaplan et al. 2012). Many studies have shown these effects — ranging from asking questions regarding the U.S. President's likability, racism prevalence, abortion stances, and foreign affairs policies (Sahin 2021; Silber et al. 2016; M a and Balyan 2019).

In the third circumstance, order effects appear to occur most when two questions are asked sequentially on the same topic or on very similar topics. This is where we could see the "judgement effect," where questions contain the same object and previous judgements may serve as standards for later comparisons (Revilla et al. 2016). More specifically, many of these datasets follow the rule of "evenhandedness," which aims to explain the famous Al Gore/Bill Clinton example (Moore 2002). In this example, a population of people is asked *Do you generally think Bill Clinton/Al Gore is honest and trustworthy?*. It is observed that respondents are more inclined to change their binary answer to the second person depending on the first person asked, such that the outcomes are "more fair." Thus, if the first question is about Clinton, people are more likely to lower their opinion of Gore to make things more "even" because Clinton generally has an overall lower score. If the first question is about Gore, people are more likely to raise their opinion of Clinton.

We take inspiration from these insights in psychology to generate two artificial datasets, which we will refer to as *D1* and *D2* in our numerical results. The first is a dataset containing order effects dependent on the general likeability scores of individuals. Given N questions of *Do you generally think this i^{th} person is honest and trustworthy?*, we assign randomly generated likeability scores $S_n \in \{0, 1\}$, to each person in question, where general public likability increases as $S_n \rightarrow 1$. Thus, when there is generally a good opinion of the person in question, the probability of someone saying *yes* to the question is high. We use a simple Markov model to produce $N!$ binary training datasets over various orders for N people in question, providing the probabilities that

a population of people answered {No, yes} in the sequence of questions. The model uses the likeability scores as hidden variable information, adjusting the outcome probabilities, by updating subsequent scores S_{n+1} dependent on the current score S_n such that it is either raised or lowered by the average difference between the two scores — thus, making outcomes "more even." The relation for which an updated score S'_{n+1} is computed is given by:

$$S'_{n+1} = S_{n+1} \pm \frac{S_{n+1} - S_n}{2} \quad (10)$$

For our second dataset, we generate a sequence of general to specific ranking questions with a corresponding random list of probabilities for public answers. If the order is changed such that a more specific question precedes a more general question, the answer to the more general question will obtain a pre-defined $x\%$ increase in the probability that the answer is *yes*. This dataset is based on authentic data regarding school bullying, where we see 45% increases of students answering *yes* to being bullied if they are asked about a specific type of bullying first (e.g., cyber bullying) (Cornell and Huang 2016).

In order to compute the strength of the order effect that exists within each dataset, we explicitly compute the LMS similarity metric (Eq. 8) over all possible data combinations that result from changing the order (note again that this is only possible due to the small scale of our experiments). For the scope of this work, it is useful to define an order effect (OE) strength that allows us to quantify the amount of order effect that exists within our dataset for our model to learn.

3 Summary of experiments

Having provided the background motivation for this investigation and the details of our model and training scheme, we now demonstrate our progress toward answering the following question: *Does the non-commutative nature of quantum measurements make quantum models constructed from sequential measurements well-suited to learn datasets generated by order effects?*

In our experiments, we run a total of 15 independent training runs for each experiment, each with different random seeds for training and data generation. The following plots report the average training loss with standard deviation error for 150 epochs. We use an Adam gradient optimizer with $stepsize = 0.1$ (Kingma and Ba 2014). Each circuit is implemented on a statevector simulator, outputting the model probability distribution directly. The experiments were implemented in the PennyLane software framework (Bergholm et al. 2022). The experiments are designed to answer two sub-questions:

1. As we increase the strength of the order effect in the dataset, does the model learn a stronger degree of non-commutativity? To address this question, we trained a $N = 2$ observable model on datasets D1 and D2 for increasing values of OE strengths, including datasets with no order effect present. For D1, the order effect (OE) strengths are controlled with likeability scores that are linearly increasing by 0.1 for the second question, where the likeability score for the first question remains fixed at 0.1 (for example, [0.1, 0.2], [0.1, 0.3], [0.1, 0.4]...). The larger the separation between scores, the larger the order effect. Note that for the likeability scores [0.1, 0.1], there is no order effect. For D2, we linearly increase the variable $x\%$ in increments of 10% from {0%, ...90%}, where 0% indicates an absence in order effect. These corresponding OE strengths for both datasets

are computed from the resulting probability distributions as described in Sect. 2.1.

We intentionally started training with two identical, and hence commuting observables, to provide a ground-zero starting point. There is no out-of-task generalization present in these experiments, as we train on samples from different orders or tasks and test on others. The training results with respect to the LMS loss and the amount of non-commutativity learned in the dataset are shown in Fig. 2. For both experiments, the model learns an increasing amount of non-commutativity throughout training, except in the absence of an order effect in the data. This suggests that the model uses non-commutativity as a resource to solve the tasks.

As shown in Fig. 2, in D1 it is very clear that the model learns a higher degree of non-commutativity as we increase

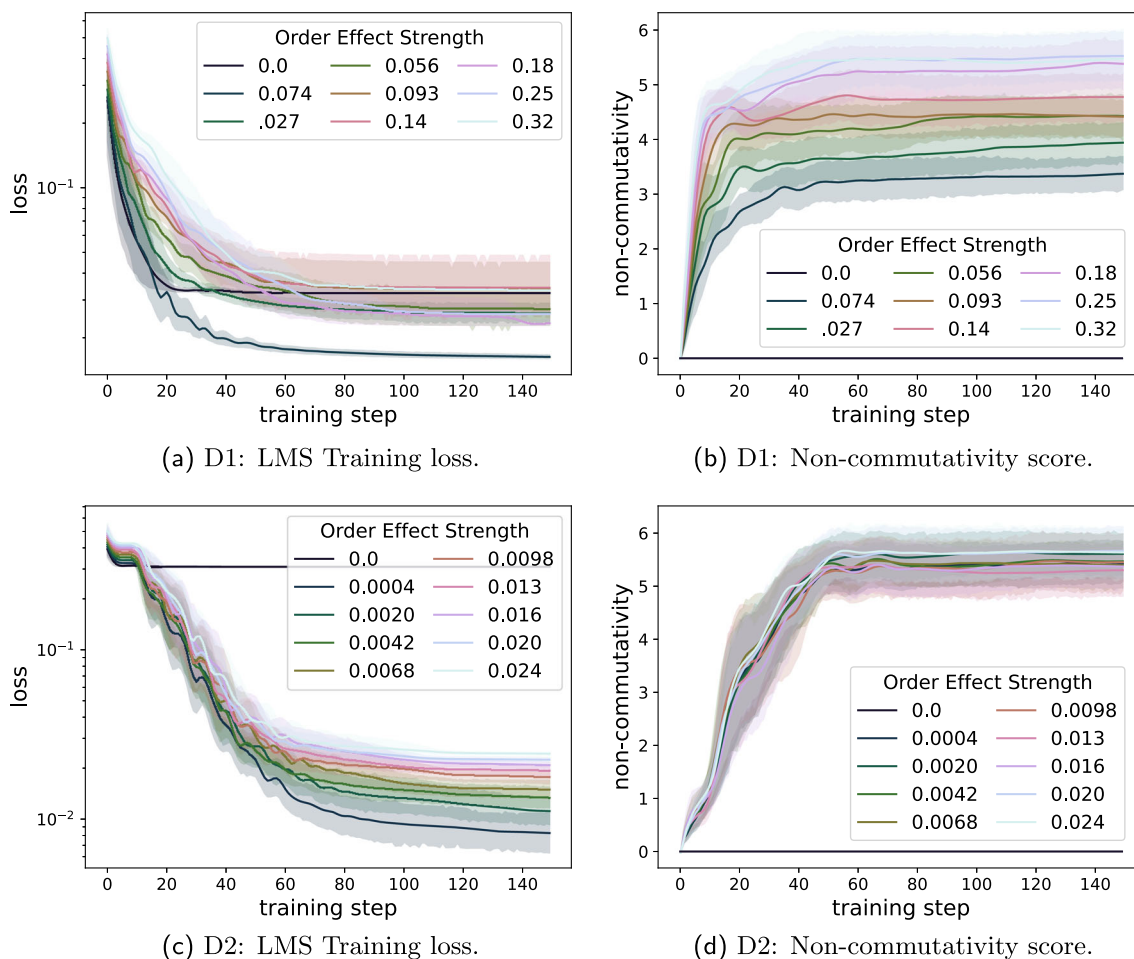


Fig. 2 LMS loss and amount of non-commutativity present in the $N = 2$ observable model throughout training for datasets D1 and D2 with various Order Effect strengths. Each figure contains results obtained over 15 independent trials with random seeds for train and test data, as well as randomness within the optimization. In **a** and **b**, we see very clearly that as one increases the strength of the order effect

in the dataset, the model learns observables with a higher degree of non-commutativity. There is no evidence of this in **c** and **d** for D2, most likely because the OE strengths are already a lot smaller than in D1. As the model is trained to learn the order effects, we see the observables learn a larger degree of non-commutativity for both D1 and D2 in all cases, except for when there is no order effect

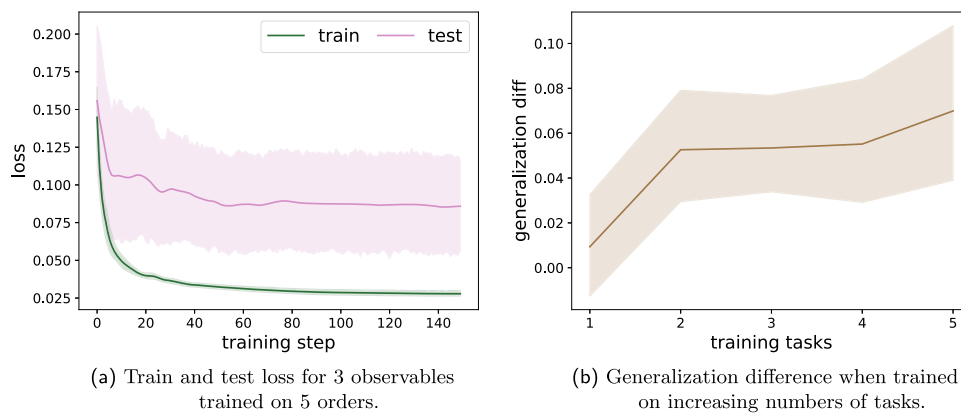


Fig. 3 Numerical results for $N = 3$ observables on D1. On the left, show an example of the decrease in train/test loss, such that one can visualize the generalization difference when trained on a specific number of tasks. On the right, we show the generalization difference with respect to training on an increased number of tasks. Each plot shows the average of 15 independent trials from randomized train/test data

and optimizer exploration. We see that the model's test loss is initially congruent with the training loss, indicating that the model is able to generalize from seen to unseen tasks even at small scale. However, the test loss does not reach the same performance as the training loss. As we train on more orders, the model is able to generalize more — which is theoretically expected behavior

the OE strength in the dataset. This is not evident in D2. We hypothesize that in D2, the OE strengths are so small (≤ 0.03) that the dataset does not have enough bias for the model to learn. We question whether OE strengths of that size are negligible for the model, and perhaps order effects in the dataset are bounded by a certain strength for the model to be able to learn them. We discuss this further in the outlook as a direction for future investigation.

For both datasets, a non-zero low strength order effect in the dataset is easiest to train on but does not correlate to as high of a non-commutativity score. In fact, there is no non-commutativity learned in the model when no order effect is present. In addition, the highest strength order effect is one of the most challenging for the model to train on, and yet, achieves the highest amount of non-commutativity. This is not necessarily surprising since distributions with greater dissimilarity pose a larger challenge for the model to find a good set of parameters. It is interesting that the model still finds a good solution with a relatively low LMS loss, and that this directly relates to the model learning a higher degree of non-commutativity. The findings support an affirmative answer to the first question, namely that the model *does use* non-commutativity to learn the order effect.

2. Can the model generalize to tasks with unseen question orders? To gain some first insights into this question, we train models with $N = \{3, 4, 5\}$ observables¹ on various

portions of tasks (orders) with dataset D1. This time, we use randomly selected likeability scores to create the tasks that comprise the test and training sets. When using a limited amount of tasks for training, we keep the number of test orders fixed at 10, except in the case where $N = 3$ and there are only 6 possible tasks. These test tasks are randomly selected over the 15 trials, along with the training tasks. We also introduce randomness into the optimization over each trial. If the model's test loss on unseen tasks decreases with its training loss on training tasks, we count this as evidence that the model is out-of-task generalizing. Below, we incorporate a metric called the *generalization difference*, which is the difference between the average LMS test loss at the beginning and end of training. A greater decrease in the test error, over the averaged set of tasks, indicates that the model is better at learning from the data.

In Fig. 3, we show the results for $N = 3$ observables when trained on various orders. Note that this is a very small-scale example with only 6 possible tasks in total, which allows us to run experiments that train on one task to all tasks *but* one. For all experiments, we see that the model's test loss on unseen tasks decreases alongside the training loss. As expected, we see that the model's generalization difference increases as we train on more orders: the model has access to more tasks, and can therefore generalize more effectively. When trained on 5 tasks and generalizing to a single task, the model's test loss on unseen tasks is closest to the training loss. However, it is also apparent that not all information about the unseen tasks is learned since training and test loss differ significantly.

When scaling our experiments to $N = 4$ and $N = 5$ observables, we do not see as clear of a trend in generalization difference as we increase the number of tasks during

¹ Note that the circuit corresponding to a 5-observable experiment uses 10 qubits on a mixed-state simulator. Larger simulations are possible with today's software but still very time-consuming, and were out of the scope of this investigation.

training. As there are more potential orders to generalize to (4 observables: 24 orders, 5 observables: 120 orders), we randomly selected 16.6%, 33.3%, and 50% of all tasks. The corresponding training and test losses are displayed in Fig. 4.

In both cases, the model learns to generalize to some extent to unseen orders, which indicates a weak *yes* to our research question. However, there is nuance in discussing the quality of the generalization performance and the amount of data required to achieve such quality. It is clear for $N = 4$ observables that as we increase the number of tasks shown to the model, the model's test error gets lower and closer to the training error, which means that generalization improves. Still, we recognize that the test error does not reach the same low loss as the training error, indicating that the out-of-task generalization performance result is weak.

For $N = 5$ observables the trend is less clear, and perhaps because the problem size is larger and hence more difficult for the model to navigate a larger parameter space. As seen in other practical ML demonstrations (Gili et al. 2023), it is possible that we are hitting a *Goldilocks region*, where increasing the amount of data shown to the model actually makes it more difficult to generalize. We also see that in the case of $N = 5$, the model is able to obtain better test errors than in the case of $N = 4$, an observation that would be interesting to study as the number of observables scales further.

Note that the experiments do not allow conclusions about the model's out-of-distribution generalization capabilities within a task, nor about whether or not the non-commutative structure reduces the amount of orders that we need to train on when scaling the problem size (which would be an interesting investigation for future work). What we can conclude from the experiments is that non-commutativity is used as a

resource, and that the model design allows it to use information gained in the training on some tasks to produce better samples from others, a sign that the inductive bias of the model may be useful for the learning task.

4 Discussion and outlook

As the field of QML is quickly expanding with a unified intention of discovering applications for quantum computers, it is not surprising that there now exists a plethora of approaches to reach this outcome. The most popular research approach of the last few years borrows language and tools from computational complexity theory to look for quantum learning settings with a built-in computational speed-up (Yoo 2014; Hinsche et al. 2021, 2022; Du et al. 2020; Liu et al. 2021; Schreiber et al. 2022; Cai et al. 2022). This *algorithmic quantum advantage* approach focuses on demonstrating that a programmable quantum device can solve a machine learning problem that no classical computer can solve in any feasible amount of time. A complimentary second approach has recently been introduced to look for quantum learning settings that contain industry relevance, where asymptotic speed-up may not be the primary measure of value, but instead the model's ability to generalize (Gili et al. 2023; Hibat-Allah et al. 2023; Gili et al. 2022a; Moussa et al. 2023). This *practical quantum advantage* approach focuses on demonstrating that a quantum algorithm or heuristic can outperform a state-of-the-art classical method for a real-world ML problem (Herrmann et al. 2023). Emphasizing quantum learning settings with trainability guarantees is an adjacent third approach in which one tries to obtain effi-

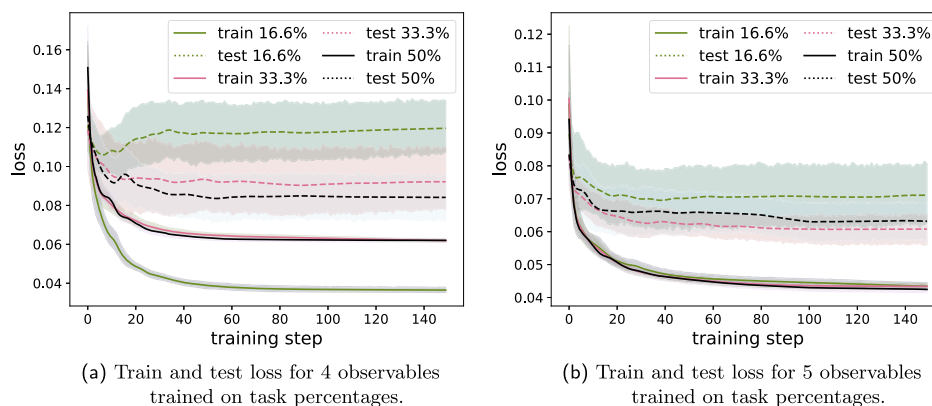


Fig. 4 Numerical results for $N = 4$ and $N = 5$ observables on D1. We show average results over 15 independent random trials with respect to the train/test data and the optimization process. In both cases, the model is able to generalize to unseen orders. It is clear for $N = 4$ observables that as we increase the number of tasks shown to the model, the model's

test error gets lower. For $N = 5$ observables this is slightly less clear, and perhaps because the problem size is larger and hence more difficult for the model to navigate a larger parameter space. However, we also see that in the case of $N = 5$, the model is able to obtain better test errors than in the case of $N = 4$

cient training methods for more general models (Abbas et al. 2023; Rudolph et al. 2022, 2023). We note that all of these approaches deliver important insights to the field of QML.

In this work, we exercise an alternative approach that focuses on finding quantum learning settings where the mathematical structure of quantum mechanics leads to an inductive bias in a quantum model that matches the structure it has to learn from a particular kind of dataset (Bowles et al. 2023; Kubler et al. 2021). This approach aims to increase our overall *understanding* of the potential natural alignments between quantum models and data structures that would lead to advantage. As for the specific scope of our work, we contribute a quantum learning setting that allows one to study an inductive bias that results from the non-commutative nature of quantum measurement. By introducing artificial datasets inspired by human psychology and a multi-task quantum generative model that contains a theoretical bias toward learning order effects, we investigate in our numerical experiments whether this bias helps the model learn datasets with increasing order effect strength and aids generalization performance. We find some initial evidence that the non-commutative bias of the model is indeed useful for learning order effects in data. We see this as a valuable insight for both the quantum cognition and QML community. Still, more evidence is necessary to gain clarity on how the resource of non-commutativity can be used in different settings. Possible questions are: *How do our findings hold up when scaling the tasks? Can we lower the out-of-task generalization error further? Does the model's non-commutative bias remain effective when the order effect in the dataset is small? Can this learning setting be adapted to isolate non-commutativity as the primary variable that leads to improved generalization performance? Does the quantum model suit other types of datasets?* Most of all, we hope this work inspires other researchers to consider alternative ways in which we can investigate the relationship between quantum mechanics and learning.

Acknowledgements K.G. would like to acknowledge all of the behind the scenes work that provided time for this research to be completed, as well as the support from her closest friends and mentors who continuously challenge her to think differently.

Author contribution K.G. All authors reviewed the manuscript.

Data availability No datasets were generated or analyzed during the current study.

Code implementation The code for this work can be found in the public repository: <https://github.com/kaitlinmgili/noncommutativity-ordereffects>

Declarations

Conflict of interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abbas A et al (2023) On quantum backpropagation, information reuse, and cheating measurement collapse. [arXiv:2305.13362](https://arxiv.org/abs/2305.13362) [quant-ph]
- Arrasmith A et al (2021) Equivalence of quantum barren plateaus to cost concentration and narrow gorges. [arXiv:2104.05868](https://arxiv.org/abs/2104.05868)
- Bao Z, Hebert M, Wang Y-X (2021) Generative modeling for multi-task visual learning. [arXiv:2106.13409](https://arxiv.org/abs/2106.13409)
- Benedetti M et al (2018) A generative modeling approach for benchmarking and training shallow quantum circuits. In: npj Quantum information, p 5
- Benedetti M et al (2019) A generative modeling approach for benchmarking and training shallow quantum circuits. In: npj Quantum information 5(1). ISSN: 2056-6387. <https://doi.org/10.1038/s41534-019-0157-8>
- Benedetti M et al (2019b). Parameterized quantum circuits as machine learning models. In: Quantum science and technology, 4(4):043001. <https://iopscience.iop.org/article/10.1088/2058-9565/ab4eb5>
- Bengio Y (2009) Learning deep architectures for AI. In: Found trends mach learn 2.1:1–127
- Bergholm V et al (2022) PennyLane: automatic differentiation of hybrid quantum-classical computations. [arXiv:1811.04968](https://arxiv.org/abs/1811.04968) [quant-ph]
- Bowles J et al (2023) Contextuality and inductive bias in quantum machine learning. [arXiv:2302.01365](https://arxiv.org/abs/2302.01365) [quant-ph]
- Busemeyer JR, Bruza PD (2012) Quantum models of cognition and decision. Cambridge University Press
- Busemeyer JR et al (2011) A quantum theoretical explanation for probability judgment errors. In: Psychological review 118.2:193
- Cai H, Ye Q, Deng, DL (2022) Sample complexity of learning parametric quantum circuits. [arXiv:2107.09078](https://arxiv.org/abs/2107.09078) [quant-ph]
- Cerezo M et al (2022) Challenges and opportunities in quantum machine learning. In: Nature computational science, 2. <https://doi.org/10.1038/s43588-022-00311-3>
- Cerezo M et al (2021) Cost function dependent barren plateaus in shallow parametrized quantum circuits. In: Nature communications, 12.1:1–12
- Cheng S, Chen J, Wang L (2018) Information perspective to probabilistic modeling: Boltzmann machines versus born machines. In: Entropy, 20.8:583. <https://doi.org/10.3390/e20080583>
- Cohen N, Shashua A (2017) Inductive bias of deep convolutional networks through pooling geometry. [arXiv:1605.06743](https://arxiv.org/abs/1605.06743) [cs.NE]
- Du Y et al (2020) Expressive power of parametrized quantum circuits. In: Physical review research, 2.3. ISSN: 2643-1564. <https://doi.org/10.1103/PhysRevResearch.2.033125>
- Feinman R, Lake BM (2018) Learning inductive biases with simple neural networks. [arXiv:1802.02745](https://arxiv.org/abs/1802.02745) [cs.CL]
- Gili K, Mauri M, Perdomo-Ortiz A (2023) Generalization metrics for practical quantum advantage in generative models. [arXiv:2201.08770](https://arxiv.org/abs/2201.08770) [cs.LG]

- Gili K, Sveistrys M, Ballance C (2022) Introducing non-linear activations into quantum generative models. <https://doi.org/10.48550/ARXIV.2205.14506>
- Gili K et al (2022a) Do quantum circuit born machines generalize?. <https://doi.org/10.48550/ARXIV.2207.13645>
- Gili K et al (2022b) Do quantum circuit born machines generalize? [arXiv:2207.13645](https://arxiv.org/abs/2207.13645). <https://arxiv.org/abs/2207.13645>
- Goyal A, Bengio Y (2022) Inductive biases for deep learning of higher-level cognition. [arXiv:2011.15091](https://arxiv.org/abs/2011.15091) [cs.LG]
- Guo Y (2016) Non-commutativity measure of quantum discord. [arXiv:1412.8325](https://arxiv.org/abs/1412.8325) [quant-ph]
- Hausler D (1988) Quantifying inductive bias: AI learning algorithms and valiant's learning framework. In: Artificial intelligence 36.2:177–221. ISSN: 0004-3702. [https://doi.org/10.1016/0004-3702\(88\)90002-1](https://doi.org/10.1016/0004-3702(88)90002-1). <https://www.sciencedirect.com/science/article/pii/0004370288900021>
- Herrmann N et al (2023) Quantum utility – definition and assessment of a practical quantum advantage. [arXiv:2303.02138](https://arxiv.org/abs/2303.02138) [quant-ph]
- Hibat-Allah M et al (2023) A framework for demonstrating practical quantum advantage: racing quantum against classical generative models. [arXiv:2303.15626](https://arxiv.org/abs/2303.15626) [quant-ph]
- Hinsche M et al (2022) A single t-gate makes distribution learning hard. <https://arxiv.org/abs/2207.03140>
- Hinsche M et al (2021) Learnability of the output distributions of local quantum circuits. [arXiv:2110.05517](https://arxiv.org/abs/2110.05517) [quant-ph]
- Holmés Z et al (2022) Connecting ansatz expressibility to gradient magnitudes and barren plateaus. In: PRX Quantum 3.1:010313
- Cornell DG, Huang FL (2016) Question order affects the measurement of bullying victimization among middle school students. In: Educ psychol meas. <https://doi.org/10.1177/0013164415622664>
- Ishibashi H, Higa K, Furukawa T (2021) Multi-task manifold learning for small sample size datasets. In: Neurocomputing 473. <https://doi.org/10.1016/j.neucom.2021.11.043>
- Kaplan S, Luchman J, Mock L (2012) General and specific question sequence effects in satisfaction surveys: integrating directional and correlational effects. In: Journal of happiness studies, 14. <https://doi.org/10.1007/s10902-012-9388-5>
- Kieferová M, Wiebe N (2017) Tomography and generative training with quantum Boltzmann machines. In: Phys Rev A, 96:062327. <https://doi.org/10.1103/PhysRevA.96.062327>. <https://link.aps.org/doi/10.1103/PhysRevA.96.062327>
- Kingma DP, Ba J (2014) Adam: a method for stochastic optimization. In: [arXiv:1412.6980](https://arxiv.org/abs/1412.6980). <https://arxiv.org/abs/1412.6980>
- Korenkevych D et al (2016) Benchmarking quantum hardware for training of fully visible boltzmann machines. In: [arXiv preprint arXiv:1611.04528](https://arxiv.org/abs/1611.04528)
- Kübler JM, Buchholz S, Schölkopf B (2021) The inductive bias of quantum kernels. [arXiv:2106.03747](https://arxiv.org/abs/2106.03747) [quant-ph]
- Leadbeater C et al (2021) F-divergences and cost function locality in generative modelling with quantum circuits. In: Entropy, 23.10. ISSN: 1099–4300. <https://doi.org/10.3390/e23101281>. <https://www.mdpi.com/1099-4300/23/10/1281>
- Liu JG, Wang L (2018a). Differentiable learning of quantum circuit born machines. In: Physical Review A, 98.6. ISSN: 2469–9934. <https://doi.org/10.1103/PhysRevA.98.062324>
- Liu JG, Wang L (2018b). Differentiable learning of quantum circuit born machines. In: PRA 98 (6 Dec. 2018) p 062324. <https://doi.org/10.1103/PhysRevA.98.062324>
- Liu Y, Arunachalam S, Temme K (2021) A rigorous and robust quantum speed-up in supervised machine learning. In: Nature physics, 17.9:1013–1017. <https://doi.org/10.1038/s41567-021-01287-z>
- M a S, Balyan P, (2019) Response order effects in online surveys: an empirical investigation. International journal of online marketing (IJOM) 4:28–44. <https://doi.org/10.4018/ijom.2014040103>
- Mcclean J et al (2018) Barren plateaus in quantum neural network training landscapes. In: Nature communications, 9
- Moore D (2002) Measuring new types of question-order effects. Public opinion quarterly - PUBLIC OPIN QUART 66:80–91. <https://doi.org/10.1086/338631>
- Moussa C et al (2023) Application of quantum-inspired generative models to small molecular datasets. [arXiv:2304.10867](https://arxiv.org/abs/2304.10867) [quant-ph]
- O'Shea K, Nash R (2015) An introduction to convolutional neural networks. [arXiv:1511.08458](https://arxiv.org/abs/1511.08458) [cs.NE]
- Perdomo-Ortiz A et al (2018) Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers. In: Quantum science and technology, 3.3:030502. <https://doi.org/10.1088/2058-9565/aab859>
- Pesah A et al (2021) Absence of barren plateaus in quantum convolutional neural networks. In: Physical review X, 11.4. <https://doi.org/10.1103/physrevx.11.041011>
- Pothos EM, Busemeyer JR (2022) Quantum cognition. Annual review of psychology 73:749–778
- Ragone M et al (2023) Representation theory for geometric quantum machine learning. [arXiv:2210.07980](https://arxiv.org/abs/2210.07980) [quant-ph]
- Revilla M, Toninelli D, Ochoa C (2016) An experiment comparing grids and item-by-item formats in web surveys completed through pcs and smartphones. In: Telematics and Informatics, 34. <https://doi.org/10.1016/j.tele.2016.04.002>
- Ruder S (2017) An overview of multi-task learning in deep neural networks. <https://doi.org/10.48550/ARXIV.1706.05098>
- Rudolph MS et al (2020) Generation of high-resolution handwritten digits with an ion-trap quantum computer. <https://doi.org/10.48550/ARXIV.2012.03924>
- Rudolph MS et al (2022) Synergy between quantum circuits and tensor networks: short-cutting the race to practical quantum advantage. In: [arXiv:2208.13673](https://arxiv.org/abs/2208.13673). <https://arxiv.org/abs/2208.13673>
- Rudolph MS et al (2023) Trainability barriers and opportunities in quantum generative modeling. [arXiv:2305.02881](https://arxiv.org/abs/2305.02881) [quant-ph]
- Şahin M (2021) Effect of item order on certain psychometric properties: a demonstration on a cyberloafing scale. In: Frontiers in psychology, 12:590545. <https://doi.org/10.3389/fpsyg.2021.590545>
- Schreiber FJ, Eisert J, Meyer JJ (2022) Classical surrogates for quantum learning models. [arXiv:2206.11740](https://arxiv.org/abs/2206.11740) [quant-ph]
- Schuld M, Killoran N (2022) Is quantum advantage the right goal for quantum machine learning? In: PRX quantum, 3.3. <https://doi.org/10.1103/prxquantum.3.030101>
- Silber H, Höhne J, Schlosser S (2016) Question order experiments in the German-European context. In: Survey methods: insights from the field. <https://doi.org/10.13094/SMIF-2016-00006>
- Stark T et al (2020) Generalization of classic question order effects across cultures. Sociological Methods and Research, p 004912411774730. <https://doi.org/10.1177/0049124117747304>
- Trueblood JS, Busemeyer JR (2011) A quantum probability account of order effects in inference. In: Cognitive science, 35.8:1518–1552
- Wang S et al (2021) Noise-induced barren plateaus in variational quantum algorithms. In: Nature communications, 12.1:1–11
- Wang Z et al (2014) Context effects produced by question orders reveal quantum nature of human judgments. In: Proceedings of the national academy of sciences of the United States of America, 111. <https://doi.org/10.1073/pnas.1407756111>
- Weinberg M et al (2018) The measurement of subjective wellbeing: Item-order effects in the personal wellbeing index—adult. In: Journal of happiness studies, 19. <https://doi.org/10.1007/s10902-016-9822-1>
- Widdows D, Rani J, Pothos EM (2023) Quantum circuit components for cognitive decision-making. In: Entropy, 25.4:548. ISSN: 1099-4300. <https://doi.org/10.3390/e25040548>
- Yearsley JM, Busemeyer JR (2016). Quantum cognition and decision theories: a tutorial. <https://doi.org/10.1016/j.jmp.2015.11.005>
- Zhang Z (2018) Improved Adam optimizer for deep neural networks. In: 2018 IEEE/ACM 26th international symposium on quality of service (IWQoS). Ieee. pp 1–2

Zhu D et al (2019) Training of quantum circuits on a hybrid quantum computer. In: Science advances, 5.10. <https://doi.org/10.1126/sciadv.aaw9918>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.