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

Kaitlin Gili[1, 2, 3], Guillermo Alonso[2], and Maria Schuld[2]

[1] University of Oxford, Oxford, United Kingdom OX1 2JD
[2] Xanadu, Toronto, ON, M5G 2C8, Canada
[3] Tufts University, Medford, MA 02155

**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.

## 1 Motivation

A critical motivation of this investigation is taken from a result that has obtained more evidence in the classical 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 [3, 5, 6]. We further define an inductive bias to be anything that encourages the model to learn specific features in the data over others. As for the inductive biases of quantum models, we can leverage non-classical mathematical structures and behaviors that exist in quantum mechanics, and search for alignments within existing datasets [2]. Here, we discuss non-commuting observable measurements in quantum mechanics and provide theoretical intuition for its structural connection to binary distributions with order effects - i.e. measured changes in the distribution due to the ordering of the variables. [10, 1]

More specifically we use the intuition behind non-commuting measurements to design a new quantum circuit model containing this intentional bias. Our proposed quantum generator leverages multi-task learning to train on multiple probability distributions of human answering patterns - artificial datasets inspired by the human cognition and psychology literature [7, 8, 9]. 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 would give support to the hypothesis that non-commutativity can create an inductive bias to learn order effects. Our results displayed in Ref. [4] 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.

## 2 The Proposed Quantum Learning Setting

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. The quantum circuit requires the number of qubits $n_{\text{qubits}}$ 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 $|\psi_0\rangle = |00\ldots0_N\rangle$, we take measurements of subsequent trainable observables, 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 [11]. The gate sequence for each observable $Q_n$ does not change when observables are permuted, and only depends on the parameters $\theta_n$. Thus, in our ansatz, the gate count and number of parameterized operations grows linearly with the number of qubits $n_{\text{qubits}}$ and the number of observables $N$. The total number of model parameters in the ansatz is given by $N(4n_{\text{qubits}} - 1)$.

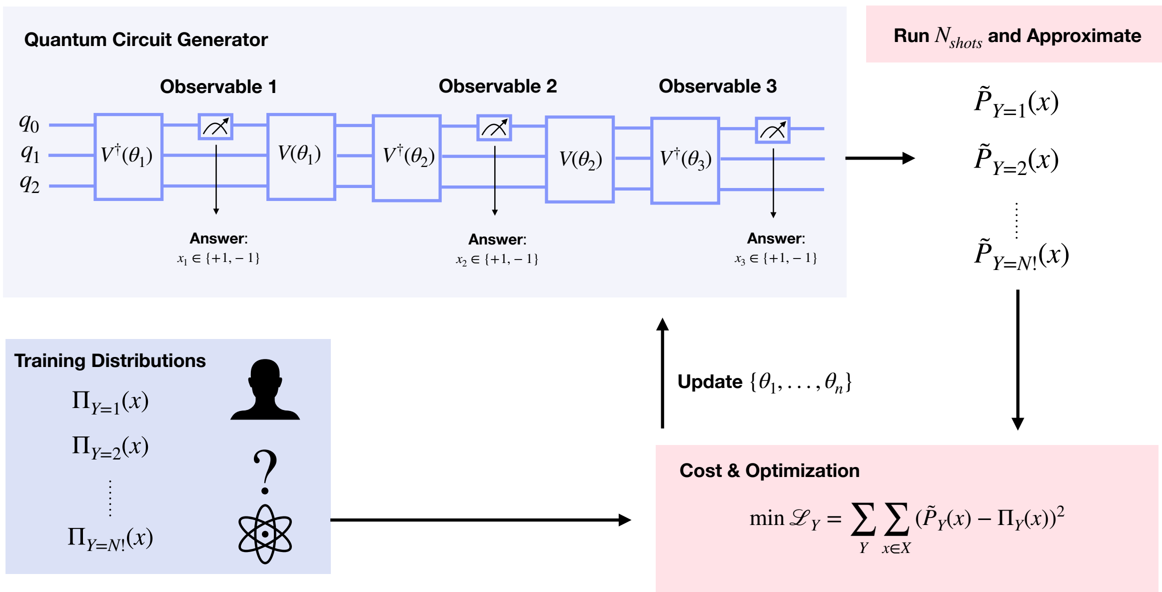


Fig. 1: **An example of the multi-task quantum generative architecture for training on binary distributions with three observables.**

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) $x_i \in \{1, -1\}$. 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 Section 4). 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 $Y$ of the observables, where $Y \in \{1, 2, ..., N!\}$. Depending on this order sequence, the circuit gives rise to one of the $N!$ probability distributions that can be used to minimize the multi-task loss function $\mathcal{L}_Y$ that computes the Least Mean Squares (LMS) 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. We provide a visualization of our entire algorithm in Figure 1.

To summarize, 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 matches 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.

## 3 Result Highlights and Discussion

By introducing artificial datasets inspired by human psychology and a multi-task quantum generative model that contains a theoretical bias towards 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. For example, we observe in Figure 2b, that as we increase the order effect in the dataset, the model learns more non-commutativity in the architecture.

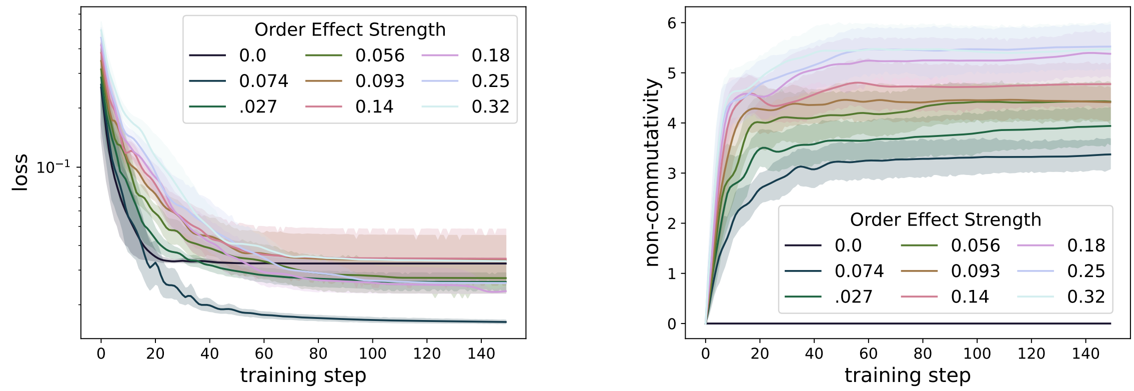


(a) D1: LMS Training loss.

(b) D1: Non-commutativity score.

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 resulted 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.

In Ref.[4], we display the rest of our initial evidence that the non-commutative bias of the model is indeed useful for learning order effects in data. We see these results 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.

## 4 Code Implementation

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

## References

[1] "A Quantum Question Order Model Supported by Empirical Tests of an A Priori and Precise Prediction". In: *Topics in cognitive science* 5 (Sept. 2013). DOI: 10.1111/tops.12040.

[2] Joseph Bowles et al. *Contextuality and inductive bias in quantum machine learning*. 2023. arXiv: 2302.01365 [quant-ph].

[3] Reuben Feinman and Brenden M. Lake. *Learning Inductive Biases with Simple Neural Networks*. 2018. arXiv: 1802.02745 [cs.CL].

[4] Kaitlin Gili, Guillermo Alonso, and Maria Schuld. *An inductive bias from quantum mechanics: learning order effects with non-commuting measurements*. 2023. arXiv: 2312.03862 [quant-ph].

[5] Anirudh Goyal and Yoshua Bengio. *Inductive Biases for Deep Learning of Higher-Level Cognition*. 2022. arXiv: 2011.15091 [cs.LG].

[6] David Haussler. "Quantifying inductive bias: AI learning algorithms and Valiant's learning framework". In: *Artificial Intelligence* 36.2 (1988), pp. 177–221. ISSN: 0004-3702. DOI: https://doi.org/10.1016/0004-3702(88)90002-1. URL: https://www.sciencedirect.com/science/article/pii/0004370288900021.

[7] Seth Kaplan, Joseph Luchman, and Landon Mock. "General and Specific Question Sequence Effects in Satisfaction Surveys: Integrating Directional and Correlational Effects". In: *Journal of Happiness Studies* 14 (Oct. 2012). DOI: 10.1007/s10902-012-9388-5.

[8] David Moore. "Measuring New Types of Question-Order Effects". In: *Public Opinion Quarterly - PUBLIC OPIN QUART* 66 (Mar. 2002), pp. 80–91. DOI: 10.1086/338631.

[9] Melissa Weinberg, Catherine Seton, and Nikki Cameron. "The Measurement of Subjective Wellbeing: Item-Order Effects in the Personal Wellbeing Index—Adult". In: *Journal of Happiness Studies* 19 (Jan. 2018). DOI: 10.1007/s10902-016-9822-1.

[10] J. M. Yearsley and Busemeyer J. R. *Quantum cognition and decision theories: A tutorial*. 2016. DOI: https://doi.org/10.1016/j.jmp.2015.11.005.

[11] Daiwei Zhu et al. "Training of quantum circuits on a hybrid quantum computer". In: *Science advances* 5.10 (2019). DOI: 10.1126/sciadv.aaw9918.