



THE UNIVERSITY
of EDINBURGH

Subspace-Preserving Quantum Machine Learning

PRESENTED BY

Léo Monbroussou

DATE

September 2nd, 2026

➤ FQC 2026 - Quantum Learning Labs, Computer
Science, UCL

Goal of this talk


The ambition of this talk is to present:

- Why computational structure is important in near term QML
- How subspace preserving structures relate to the main QML challenges
- What are the open questions and hope for the future of this field

Subspace preserving Variational Quantum Algorithms


Variational quantum algorithms

The goal is to approximate a function $f^* : x \mapsto y$ from a dataset $\{x_i, y_i\}_{i=1}^D$:





Variational quantum algorithms

The goal is to approximate a function $f^* : x \mapsto y$ from a dataset $\{x_i, y_i\}_{i=1}^D$:




Variational quantum algorithms

The goal is to approximate a function $f^* : x \mapsto y$ from a dataset $\{x_i, y_i\}_{i=1}^D$:




Variational quantum algorithms

The goal is to approximate a function $f^* : x \mapsto y$ from a dataset $\{x_i, y_i\}_{i=1}^D$:




How to design near term QML algorithms ?




How to design near term QML algorithms ?


The neural network should
converge to a solution



How to design near term QML algorithms ?




How to design near term QML algorithms ?





Subspaces and Lie algebras

Subspace preserving algorithms can be created using a non universal set of hamiltonian: $\{\textcolor{red}{H}_1, \textcolor{teal}{H}_2, \textcolor{lightgreen}{H}_3, \textcolor{yellow}{H}_4, \dots\}$





Subspaces and Lie algebras

Subspace preserving algorithms can be created using a non universal set of hamiltonian: $\{\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3, \mathbf{H}_4, \dots\}$



Subspaces and Lie algebras


Subspace preserving algorithms can be created using a non universal set of hamiltonian: $\{\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3, \mathbf{H}_4, \dots\}$



$$\mathfrak{g} = \text{span} \langle iH_1, iH_2, iH_3, iH_4, \dots \rangle_{Lie}$$


Subspaces and Lie algebras

Subspace preserving algorithms can be created using a non universal set of hamiltonian: $\{\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3, \mathbf{H}_4, \dots\}$



Symmetries and subspaces

The use of symmetries in the choice of the gates allows to separate subspaces: $\mathfrak{g} \simeq \bigoplus_{k=0}^n \mathfrak{g}_k$




Hamming-weight preserving quantum circuits

$$RBS(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) & 0 \\ 0 & -\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

*The Reconfigurable BeamSplitter
Gate is Hamming Weight Preserving.*

Considering 4 qubits, the computational basis can be rearranged according to the Hamming-weight:




Hamming-weight preserving quantum circuits

$$RBS(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) & 0 \\ 0 & -\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

*The Reconfigurable BeamSplitter
Gate is Hamming Weight Preserving.*

Considering 4 qubits, the computational basis can be rearranged according to the Hamming-weight:




S. Johri, et al., "Nearest centroid classification on a trapped ion quantum computer," *np Quantum Information*, Aug. 2021.

Hamming-weight preserving quantum circuits

$$RBS(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) & 0 \\ 0 & -\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

*The Reconfigurable BeamSplitter
Gate is Hamming Weight Preserving.*


Considering 4 qubits, the computational basis can be rearranged according to the Hamming-weight:



S. Johri, et al., "Nearest centroid classification on a trapped ion quantum computer," *np Quantum Information*, Aug. 2021.


Controllable Lie algebras

A dynamical Lie algebra is controllable if it is simple and full rank. The dynamical Lie group gives a precise description of the achievable unitary matrices:



Controllability at different stages:

Quantum signals are bounded by the controllability in the different mathematical spaces of interest:





By imposing a specific structure to the computation, we can expect the output model to have particular properties regarding its expressivity, training, and ability to avoid dequantization.

Trainability of subspace preserving VQCs

Expressivity versus Trainability



The distribution of the achievable unitary matrices is an important figure of merit for expressivity.



Z. Holmes, et al., "Connecting Ansatz Expressibility to Gradient Magnitudes and Barren Plateaus," *PRX Quantum*, Jan. 2022.

Expressivity versus Trainability


The distribution of the achievable unitary matrices is an important figure of merit for expressivity.




Z. Holmes, et al., “Connecting Ansatz Expressibility to Gradient Magnitudes and Barren Plateaus,” *PRX Quantum*, Jan. 2022.

Expressivity versus Trainability

The distribution of the achievable unitary matrices is an important figure of merit for expressivity.



Vanishing Gradient Phenomena: **Barren Plateau**



J. R. McClean, et al., "Barren plateaus in quantum neural network training landscapes," [Nature Communications](#), Nov. 2018.

M. Larocca, et al., "Barren plateaus in variational quantum computing," [Nature Reviews Physics](#), Apr. 2025.

Vanishing gradient and subspaces

By considering **Haar random unitary matrices on the subspace**:

$$U(\theta) = \begin{pmatrix} 1 & & & & 0 \\ & \ddots & & & \\ & & U^1(\theta) & & \\ & & & U^2(\theta) & \\ & & & & \ddots \\ \vdots & & & & & U^{n/2}(\theta) & \\ & & & & & & \vdots \\ & & & & & & & U^{n-2}(\theta) \\ & & & & & & & & U^{n-1}(\theta) \\ 0 & & & & & & & & & 1 \end{pmatrix}$$

d_1 d_2 $d_{n/2}$ d_{n-2} d_{n-1}


$\mathbb{E}_{U \sim \mathcal{U}(SU(d))} [f_U^2] = \Omega(\text{poly}(n)^{-1})$

$\mathbb{E}_{U \sim \mathcal{U}(SU(d))} [f_U^2] = O(e^{-n})$

M. Larocca, et al., "Barren plateaus in variational quantum computing," *Nature Reviews Physics*, Apr. 2025.


Vanishing gradient and subspaces

Most of the results on the trainability of quantum neural networks are based on the distribution of the achievable unitary matrices, which is different from the dimensionality of the set of achievable unitary matrices (bounded by the lie algebra dimension).




Vanishing gradient and Lie algebras

By considering **Haar randomness** over some the Lie group:



Quantum model:
 $f_U(\rho, O)$

Concentration diagnosis



$$\mathbb{E}_U[f_U^2] = \sum_{k=0}^M \frac{\|P_k(\rho)\|_2^2 \|P_k(O)\|_2^2}{d_k^{(n)}}$$

The variance of the quantum model depends on how much the input state and the observable are part of each Lie algebras.


M. Ragone, et al., "A Lie algebraic theory of barren plateaus for deep parameterized quantum circuits," [Nature Communications](#), Aug. 2024.

E. Fontana, et al., "Characterizing barren plateaus in quantum ansätze with the adjoint representation," [Nature Communications](#), Aug. 2024.

Classical Simulation versus Trainability

Classical simulation is possible when considering input states and observable part of polynomially large Lie algebras:

Concentration diagnosis



$$\mathbb{E}_U[f_U^2] = \sum_{k=0}^M \frac{\|P_k(\rho)\|_2^2 \|P_k(O)\|_2^2}{d_k^{(n)}}$$

Lie-algebraic classical simulations for quantum computing

Matthew L. Goh^{1,2,*}, Martin Larocca^{1,3}, Lukasz Cincio¹, M. Cerezo⁴ and Frédéric Sauvage^{1,5}

¹Theoretical Division, [Los Alamos National Laboratory](#), Los Alamos, New Mexico 87545, USA

²Department of Materials, [University of Oxford](#), Parks Road, Oxford OX1 3PH, United Kingdom

³Center for Nonlinear Studies, [Los Alamos National Laboratory](#), Los Alamos, New Mexico 87545, USA

⁴Information Sciences, [Los Alamos National Laboratory](#), Los Alamos, New Mexico 87545, USA

⁵Quantinuum, Partnership House, Carlisle Place, London SW1P 1BX, United Kingdom

Enabling Lie-Algebraic Classical Simulation beyond Free Fermions

Adelina Bärligea^{1,*}, Matthew L. Sims-Goh^{2,3} and Jakob S. Kottmann^{1,4}

¹Institute for Computer Science, [University of Augsburg](#), 86159 Augsburg, Germany

²Quantum Motion Technologies Pty Ltd, Eveleigh NSW 2015, Australia

³School of Physics, [The University of Melbourne](#), Parkville VIC 3010, Australia

⁴Center for Advanced Analytics and Predictive Sciences,
[University of Augsburg](#), 86159 Augsburg, Germany

M. Cerezo, et al., “Does provable absence of barren plateaus imply classical simulability?”, [Nature Communications](#), Aug. 2024.

**Why using structure and
subspace preserving
algorithms?**


Why structure is important

Subspace preserving circuits allow to:

- Mimic classical ML architectures and ensure good performances at scale.
- Reach polynomial advantage with lower hardware constraints, and thus potential utility.
- Go beyond what theory gives us.

Expressivity and structure

Structure is very important in classical machine learning. Many QML algorithms are based on symmetries and mimic the behavior of classical machine learning neural networks:



- E. A. Cherrat, I. Kerenidis, N. Mathur, J. Landman, M. Strahm, and Y. Y. Li, “Quantum Vision Transformers”, *Quantum*, vol. 8, p. 1265, Feb. 2024.


- > Transformer architectures

- S. Raj, B. Coyle, “QuIC: Quantum-Inspired Compound Adapters for Parameter Efficient Fine-Tuning”, *arXiv:2502.06916v2*, Feb. 2025.

- > Quantum Inspired Adapters for LLM architectures

Convolutional neural networks


CNNs are used for image classification, video analysis, natural language processing, anomaly detection, drug discovery, ...



Y. Lecun, et al., "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, Nov. 1998.

Subspace preserving encoding


The subspace preserving convolutional layer is based on the *tensor encoding*, that allows to preserve the image structure.



Subspace preserving quantum convolutional neural network architectures, L. Monbroussou, Jonas Landman, Letao Wang, Alex Grilo, and Elham Kashefi, [Quantum Science and Technology](#), March 2025.

Measurements based operations

The pooling layers are measurement based, non linear operations that still need to preserve the structure of the images so they can be alternate with



QCNN polynomial advantage


Is polynomial advantage and training guarantees enough for quantum utility ?

	Convolutional	Pooling	Orthogonal dense	Dense
Classical Time Complexities	$\mathcal{O}(\prod_{i=1}^k K_i^2 \cdot d_i^2)$	$\mathcal{O}(\prod_{i=1}^k d_i)$	$\mathcal{O}(p \cdot \binom{n}{k})$	$\mathcal{O}((\sum_{i=1}^k d_i)^2)$
Quantum Layer Depth	$\mathcal{O}(\log(K))$	$\mathcal{O}(1)$	$\mathcal{O}(\frac{p}{n})$	—

Subspace preserving quantum convolutional neural network architectures, L. Monbroussou, Jonas Landman, Letao Wang, Alex Grilo, and Elham Kashefi, [Quantum Science and Technology](#), March 2025.

Linear optical circuits


Circuits only made of beam splitters and phase shifters, that preserve the number of particles are subspace preserving. **These circuits can offer high repetition rates, especially when considering low number of particles.**




Linear optical circuits

Circuits that preserve the number of particles are subspace preserving.

Circuits made of linear optics are limited in their controllability by the **photonic homomorphism**.



$$W(\Theta) =$$



S. Aaronson and A. Arkhipov, "The computational complexity of linear optics," in *Proceedings of the forty-third annual ACM symposium on Theory of computing, STOC'11, (New York, NY, USA), pp. 333–342, Association for Computing Machinery, 2011.*

Photonic QCNN



Photonic quantum convolutional neural networks with adaptive state injection

Léo Monbroussou^{a,b}, Beatrice Polacchi^c, Verena Yacoub^a, Eugenio Caruccio^c, Giovanni Rodari^{b,c}, Francesco Hoch^{b,c}, Gonzalo Carvacho^{b,c}, Nicolò Spagnolo^{b,c}, Taira Giordani^{b,c}, Mattia Bossi^{b,d,e}, Abhiram Rajan^{d,e}, Niki Di Giano^{d,e}, Riccardo Albiero^{b,e}, Francesco Ceccarelli^{b,e}, Roberto Osellame^{b,e,*}, Elham Kashefi^{b,a,f,*} and Fabio Sciarrino^{b,c,*}

^aSorbonne Université, Laboratoire d'Informatique de Paris 6, CNRS, Paris, France

^bCEMIS, Direction Technique, Naval Group, Ollioules, France

^cSapienza Università di Roma, Dipartimento di Fisica, Roma, Italy

^dPolitecnico di Milano, Dipartimento di Fisica, Milano, Italy

^eConsiglio Nazionale delle Ricerche, Istituto di Fotonica e Nanotecnologie (IFN-CNR), Milano, Italy

^fUniversity of Edinburgh, School of Informatics, Edinburgh, United Kingdom

Highlighted
Authors:



Verena Yacoub

LIP6



Beatrice Polacchi

La Sapienza



Hugo Thomas

LIP6, INRIA, Quandela



Elliott Z. Mamon

LIP6

PHYSICAL REVIEW RESEARCH 7, 033051 (2025)

Toward quantum advantage with photonic state injection

Léo Monbroussou^{a,1,2,*}, Eliott Z. Mamon^{b,1,†}, Hugo Thomas^{b,3,1,4,†}, Verena Yacoub^{a,1,†},
Ulysse Chabaud^{b,4} and Elham Kashefi^{1,5}

¹*Sorbonne Université, CNRS, LIP6, 75005 Paris, France*

²*CEMIS, Direction Technique, Naval Group, 83190 Ollioules, France*

³*Quandela, 7 rue Léonard de Vinci, 91300 Massy, France*

⁴*DIENS, Ecole Normale Supérieure, PSL University, CNRS, INRIA, 45 rue d'Ulm, 75005 Paris, France*


⁵*School of Informatics, University of Edinburgh, 10 Crichton Street, EH8 9AB Edinburgh, United Kingdom*



(Received 21 October 2024; accepted 6 April 2025; published 11 July 2025)


Photonic QCNN

We can adapt the subspace preserving QCNN for photonic platform using state-injection:



Hardware-aware photonic models

We can adapt the Hamming-weight preserving QCNN for photonic platform using passive linear optical platforms and a measurement based operation named state-injection:




Subspace preserving quantum convolutional neural network architectures, L. Monbroussou, et al., *Quantum Science and Technology*, March 2025.

Photonic Quantum Convolutional Neural Networks with Adaptive State Injection, L. Monbroussou, et al., *Advanced Photonics*, November 2025.


Graph neural networks

The GNN success rests on **message passing**, and are enhanced by properties such as **permutation equivariance**, or **Weisfeiler-Leman (WL) capacity**.




Graph neural networks

The GNN success rests on **message passing**, and are enhanced by properties such as **permutation equivariance**, or **Weisfeiler-Leman (WL) capacity**.





Graph neural networks

Good performances on different datasets, with a small number of parameters. Large subspaces are used to encode the data and ensure the theoretical properties, potentially exponentially hard to simulate, but no trainability guarantee.



GNNs training: using pre-training

As for classical GNNs, our architecture can be pre-trained on small graph instances where the circuit is classical stimuable and post-trained on large graph instances. **We numerically show that our architecture avoid trainability issues.**



Quantum graph neural network



Quantum Physics

[Submitted on 25 Jun 2026]

Scalable Message-Passing Quantum Graph Neural Networks in the Weisfeiler-Leman Hierarchy

Snehal Raj, Brian Coyle, Léo Monbroussou, André J. Ferreira-Martins, Renato M. S. Farias, Elham Kashefi

Highlighted
Authors:




Snehal Raj
LIP6 - QC Ware



Brian Coyle
QSL - Fujitsu

Passive linear optical models

When studying the passive linear optical models simulation and trainability, we noticed that we can find trainable models that cannot be simulated with an exponentially small error.



H. Mhiri, H. Thomas, L. Monbroussou, U. Chabaud, Z. Holmes, E. Kashefi "Boson sampling beyond the dilution regime: second moments and anti-concentration," [arXiv:2604.14323](#), April. 2026.


L. Monbroussou, H. Thomas, H. Mhiri, Z. Holmes, E. Kashefi "Classical simulation and model concentration in passive linear optics," [arXiv:2607.24728](#), July. 2026.

Passive linear optical models

When studying the passive linear optical models simulation and trainability, we noticed that we can find trainable models that cannot be simulate with an exponentially small error.

$$\mathbb{E}_U[f_U^2] = \sum_{k=0}^n \frac{\|P_k^{(n)}(\rho)\|_2^2 \|P_k^{(n)}(O)\|_2^2}{d_k^{(n)}}$$

Input State
irrep projection:
 $\{\|P_k^{(n)}(\rho)\|_2^2\}_{k \leq n}$
 $\sum_{k=0}^n \|P_k^{(n)}(\rho)\|_2^2 = \text{Tr}[\rho^2]$



Observable
irrep projection:
 $\{\|P_k^{(n)}(O)\|_2^2\}_{k \leq n}$
 $\sum_{k=0}^n \|P_k^{(n)}(O)\|_2^2 = \|O\|_2^2$

Irrep dimension: $\{d_k^{(n)}\}_{k \leq n}$

H. Mhiri, H. Thomas, L. Monbroussou, U. Chabaud, Z. Holmes, E. Kashefi “Boson sampling beyond the dilution regime: second moments and anti-concentration,” [arXiv:2604.14323](#), April. 2026.

L. Monbroussou, H. Thomas, H. Mhiri, Z. Holmes, E. Kashefi “Classical simulation and model concentration in passive linear optics,” [arXiv:2607.24728](#), July. 2026.

Any questions?

➤ Imonbrou@ed.ac.uk