

Symmetry Is Not a Shortcut: Complexity and Learning Barriers

in Orthogonal and Symplectic Quantum Circuits

Oxana Shaya^{1*} Zoë Holmes^{2,3} Christoph Hirche¹ Armando Angrisani^{2,3†}

¹Institute for Information Processing (tnt/L3S), Leibniz Universität Hannover, Germany

²Institute of Physics, Ecole Polytechnique Fédérale de Lausanne, Switzerland

³Centre for Quantum Science and Engineering, Ecole Polytechnique Fédérale de Lausanne, Switzerland

Abstract. We study whether symmetry-constrained quantum circuits can avoid the complexity and learning barriers usually associated with Haar-random models. Focusing on circuits drawn from the special orthogonal and unitary symplectic groups, we show that these structured ensembles still generate states of exponentially large strong state complexity, with analogous behavior for group designs. We further prove average-case hardness for statistical-query learning of their output distributions. This has direct implications for quantum circuit Born machines: restricting an overly expressive generative ansatz to a structured subgroup does not, by itself, restore learnability or trainability. Our results clarify both the usefulness and the limitations of symmetry as a design principle for quantum learning and generative modeling.

Keywords: quantum circuit Born machines, statistical-query learning, random quantum circuits, orthogonal and symplectic groups, strong state complexity.

1 Motivation

Random quantum circuits are a basic model of generic quantum behavior. They are used to study scrambling, quantum chaos, benchmarking, tomography, random-circuit sampling, and the limits of quantum learning. They also underlie quantum circuit Born machines (QCBMs), where a parametrized quantum circuit defines a generative model through the probability distribution obtained by measuring its output state. In such models one typically has sample access rather than exact probabilities, so training objectives and gradient estimates are built from empirical averages. This makes learnability and trainability naturally connected to statistical-query access.

A common intuition in quantum machine learning is that structure should help. If an ansatz is restricted to a symmetry-respecting family, such as real-valued or symplectic circuits, then perhaps it becomes more controllable while retaining useful expressive power. This intuition is important: symmetry is used to encode inductive bias, reduce the search space, and adapt circuit families to data, physics, or hardware constraints. But it also raises a basic question. Does symmetry really remove the complexity barriers that appear for random quantum models, or can structured circuit families remain just as hard to understand and learn?

We test this question for the classical compact groups acting on $D = 2^n$ dimensions: the special unitary group $SU(D)$, the special orthogonal group $SO(D)$, and the unitary symplectic group $Sp(D/2)$. The special unitary group is the familiar Haar-random benchmark. The new point is that the real and symplectic ensembles, although highly structured, exhibit the same qualitative complexity behavior in the regimes we analyze.

Our main message is that symmetry is not, by itself, a shortcut to simplicity. Orthogonal and symplectic randomness already generates states that cannot be recognized by small circuits,

* shayaoxana@gmail.com

† armando.angrisani@epfl.ch

produces enormous families of well-separated high-complexity states, and yields Born distributions that are average-case hard to learn from noisy aggregate information. This makes these ensembles useful models for quantum information, while also identifying a limitation of symmetry-based ansatz design for QCBMs and related quantum generative models.

2 Results

Typical states are strongly complex. We use the strong state complexity of Brandão, Chennissany, Hunter-Jones, Kueng, and Preskill. Operationally, a pure state $|\psi\rangle$ has large strong complexity if no measurement implementable by a small circuit can reliably distinguish $|\psi\rangle\langle\psi|$ from the maximally mixed state I/D . This definition is useful because it speaks directly about what a bounded quantum experiment can certify, rather than about a particular preparation circuit.

For a random state $|\psi\rangle = U|0^n\rangle$ with U drawn from any of $SU(D)$, $SO(D)$, or $Sp(D/2)$, low strong complexity is overwhelmingly unlikely. More concretely, for any fixed distinguishing advantage, the probability that $|\psi\rangle$ can be distinguished from I/D by a measurement of size r is exponentially small in D until

$$r \simeq D/\log n = 2^n/\log n.$$

Thus typical states from the real and symplectic ensembles are exponentially complex in the same operational sense as typical Haar-unitary states.

We also extend the analysis to approximate designs over these groups. Designs are finite or efficiently generable ensembles that match Haar moments up to some order. The design result gives an explicit tradeoff between design order, approximation accuracy, and the size of circuits that fail to distinguish the state from the maximally mixed state. This is important for near-term and architecture-constrained settings, where one often implements an approximate design rather than exact Haar sampling.

High-complexity states are plentiful and well separated. The complexity phenomenon is not caused by a few exceptional states. We prove a packing statement: inside each Haar-induced ensemble there are

$$\exp(\Omega(D)) = 2^{2^{\Omega(n)}}$$

pure states that all have large strong state complexity and are pairwise almost maximally far apart in trace distance. Intuitively, the structured ensembles contain an enormous landscape of mutually distinct hard-to-recognize states. For approximate designs, we derive an explicit lower bound on the number of such separated states in terms of the design parameters and the desired separation threshold.

This geometric statement strengthens the first result. It says not only that a typical draw is hard for small distinguishers, but also that high-complexity states occupy the ensemble in a robust and well-spread way. This is the kind of geometry expected from truly random quantum states, and it persists under orthogonal and symplectic structure.

Born distributions are hard to learn. The result most directly connected to QCBMs concerns the classical distribution obtained by measuring $U|0^n\rangle$ in the computational basis. The learner is given statistical-query access: it can ask for approximate expectation values of bounded functions of the measurement outcome. This captures the finite-shot situation in which a statistic can only be estimated up to sampling noise, and it matches the form of many objectives used to train quantum generative models from samples.

A key ingredient is that the output distributions from these ensembles are typically far from uniform. We compute the relevant average total-variation scale using a Gaussian integration method. For the real ensemble the leading constant is $\sqrt{2/(\pi e)}$, while for the unitary and symplectic ensembles it is $1/e$, up to corrections vanishing as $D^{-1/2}$. Combining this non-uniformity with concentration of measure and a general statistical-query lower-bound framework, we prove that learning a prescribed fraction of the Born distributions requires doubly exponentially many statistical queries

in natural parameter regimes:

$$q = 2^{2^{\Omega(n)}}.$$

This has two complementary consequences. On the positive side, these far-from-uniform output distributions are relevant for verification tasks such as heavy-output generation, where nontrivial sampling structure is the resource. On the cautionary side, they show that restricting an overly expressive QCBM ansatz to $\text{SO}(D)$ or $\text{Sp}(D/2)$ does not, by itself, restore learnability from finite-shot statistics or remove the associated trainability barrier. Symmetry can be useful, but it must be combined with additional control of expressivity, initialization, data dependence, or optimization dynamics.

3 Main ideas

The proofs use three complementary ideas. First, concentration of measure on compact groups shows that a fixed small measurement has almost the same acceptance probability on a random structured state as on the maximally mixed state. A union bound over all measurements implementable by circuits of size r then turns this fixed-measurement statement into a strong-complexity statement. The counting of circuits is elementary; the nontrivial input is that Lévy-type concentration is strong enough on $\text{SO}(D)$ and $\text{Sp}(D/2)$, just as it is for the full unitary group.

Second, the approximate-design extension replaces exact Haar concentration with moment information. Large-deviation bounds for designs allow the same proof strategy to survive when the ensemble only matches the first few moments of Haar randomness. This gives quantitative finite-design statements rather than merely asymptotic intuition.

Third, for the learning lower bound, we avoid heavy Weingarten calculations by reducing averages over the relevant random states to averages over normalized Gaussian vectors. The real, complex, and symplectic cases then become parallel calculations. These estimates identify how far a typical Born distribution is from uniform and how sharply it concentrates. Plugging these estimates into the statistical-query lower bound yields the doubly exponential query requirement.

4 Impact and outlook

The work clarifies which aspects of random-unitary complexity depend on the full unitary group and which do not. The answer is useful for quantum information: many phenomena normally attributed to Haar randomness survive in structured, physically meaningful ensembles. Orthogonal and symplectic circuits can therefore serve as symmetry-aware models of random quantum behavior for pseudorandomness, chaos-inspired modeling, benchmarking, and sampling-based verification.

The same results also clarify a limitation for quantum machine learning. Symmetry is an important design principle, but our lower bounds show that it is not a cure for overexpressivity. If a QCBM or related generative ansatz remains too random, then finite-shot access to empirical averages may still be insufficient for learning a typical output distribution. This supports a more nuanced design principle: useful structure should not merely restrict the group, but should also shape the reachable distributions and the training trajectory.

Several directions remain open. One is to study how quickly complexity grows in concrete orthogonal or symplectic gate architectures at finite depth. Another is to understand how noise interacts with structured randomness; unital noise may enhance concentration, while nonunital noise may change the picture. A third is to identify which additional restrictions make symmetry-aware quantum models genuinely learnable. The present results provide a baseline: orthogonal and symplectic structure is already rich enough to support Haar-like complexity, so learnable QCBMs require more than symmetry alone.

A technical version of this work is available as [arXiv:2509.07573](https://arxiv.org/abs/2509.07573).

References

- [1] O. Shaya, Z. Holmes, C. Hirche, and A. Angrisani. On the complexity of quantum states and circuits from the orthogonal and symplectic groups. [arXiv:2509.07573](https://arxiv.org/abs/2509.07573), 2025.

- [2] F. G. S. L. Brandão, W. Chemissany, N. Hunter-Jones, R. Kueng, and J. Preskill. Models of quantum complexity growth. *PRX Quantum*, 2:030316, 2021.
- [3] M. Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM*, 45(6):983–1006, 1998.
- [4] A. Nietner, M. Ioannou, R. Sweke, R. Kueng, J. Eisert, M. Hinsche, and J. Haferkamp. On the average-case complexity of learning output distributions of quantum circuits. *arXiv:2305.05765*, 2023.
- [5] M. Benedetti, D. Garcia-Pintos, Y. Nam, and A. Perdomo-Ortiz. A generative modeling approach for benchmarking and training shallow quantum circuits. *npj Quantum Information*, 5:45, 2019.
- [6] A. K. Hashagen, S. T. Flammia, D. Gross, and J. J. Wallman. Real randomized benchmarking. *Quantum*, 2:85, 2018.
- [7] Y. Nakata, D. Zhao, T. Okuda, E. Bannai, Y. Suzuki, S. Tamiya, K. Heya, Z. Yan, K. Zuo, S. Tamate, Y. Tabuchi, and Y. Nakamura. Quantum circuits for exact unitary t -designs and applications to higher-order randomized benchmarking. *PRX Quantum*, 2:030339, 2021.
- [8] L. Grevink, J. Haferkamp, M. Heinrich, J. Helsen, M. Hinsche, T. Schuster, and Z. Zimborás. Will it glue? On short-depth designs beyond the unitary group. *arXiv:2506.23925*, 2025.
- [9] J. Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.