

Quantum spectral proper orthogonal decomposition

Jona Erle*

David Halpenny†

Gianmarco Mengaldo‡

Alessandro Luongo§

Abstract. Spectral proper orthogonal decomposition (SPOD) is a widely used method for extracting physically meaningful dominant structures from turbulent flows, at a classical cost $\tilde{O}(MN_f L^2)$ dominated by a spectral decomposition at each of N_f frequencies. We propose a quantum algorithm, qSPOD, that applies a quantum Fourier transform to represent these frequency-wise problems as a single block-diagonal operator whose eigenvectors are self-labeled by their frequency. Consequently, a single quantum PCA routine on this block-diagonal operator can recover the dominant SPOD modes across all frequencies at once, resulting in an exponential improvement in the data dimensions M, L , and N_f under standard quantum state-preparation assumptions. Combining quantum-inspired sample-and-query techniques with sparse FFTs, we show that this advantage can be recovered classically up to large polynomial overheads, forfeiting any exponential quantum advantage. SPOD thus provides a concrete, physically motivated setting where quantum speedups, Fourier structure, and dequantization interact nontrivially.

1 Motivation

Spectral proper orthogonal decomposition (SPOD) is used to extract energy-ranked, frequency-resolved coherent structures from statistically stationary turbulent flows. In its original form, the decomposition dates back to Lumley [1, 2], while the modern formulation is due to Towne, Schmidt, and Colonius [3, 4]. Since every SPOD mode oscillates at a single frequency—unlike space-only POD, whose modes generally mix dynamics across many frequencies—SPOD is well-suited for identifying physically meaningful structures in the flow and is used across a wide range of applications [5, 6].

From an algorithmic point of view, SPOD operates as follows: Given M real-valued variables sampled at $N_t = N_f L$ time steps, one segments the data into L blocks of length N_f . We call the resulting data structure $Q \in \mathbb{R}^{M \times L \times N_f}$ and denote its entries as Q_{mlt} . A temporal DFT on each block yields for every frequency bin $k \in \{0, \dots, N_f - 1\}$ a coefficient matrix $\hat{Q}_k \in \mathbb{C}^{M \times L}$ whose entry (m, l) is the DFT coefficient of variable m in block l at frequency k . Under the assumption that the data-generating process is statistically stationary and ergodic, one can treat each block l as an independent realization of a statistical ensemble and approximate the cross-spectral density as $C_k = \hat{Q}_k \hat{Q}_k^\dagger \in \mathbb{C}^{M \times M}$. Diagonalizing each C_k gives non-zero eigenpairs $(\lambda_{k,r}, \psi_{k,r})$, where $r \in \{1, \dots, R\}$, $R := \min(M, L)$. Each $\psi_{k,r}$ is a space-time coherent mode, i.e. a spatial structure evolving at the single frequency f_k , with energy $\lambda_{k,r}$. While SPOD outputs every eigenpair, in practice, one is often only interested in the globally dominant structures, i.e. the pairs (k, r) of largest $\lambda_{k,r}$ across all k and r . These dominant modes are then either used to visualize the coherent structures or to compute scalar diagnostics of them, e.g. its energy $\lambda_{k,r}$, its frequency k , or its overlap $\langle a, \psi_{k,r} \rangle$ to a test field a .

The cost of SPOD thus splits into the $M \times L$ DFTs and the N_f per-frequency spectral decompositions, resulting in a total runtime of

$$\underbrace{O(MLN_f \log N_f)}_{\text{temporal DFTs}} + \underbrace{O(N_f ML^2)}_{\text{per-frequency SVD}} = \tilde{O}(MN_f L^2), \quad (1)$$

where we assume $L \leq M$ (which is generally true in practice) and obtain each spectral decomposition from the thin SVD of $\hat{Q}_k$, so that the spectral stage costs $O(ML^2)$ rather than $O(M^2 L + M^3)$ per frequency. The second term dominates whenever $L > \log N_f$, so the per-frequency PCA sets the overall cost.

While SPOD is classically efficient, that is polynomial in M, N_f , and L , two observations invite a quantum treatment. First, the quantum Fourier transform realizes the same DFT on an amplitude-encoded signal with exponentially fewer gates, i.e. $O(\text{poly log } N_f)$ versus the FFT's $O(N_f \log N_f)$ [7]. Second, quantum principle component analysis (qPCA) extracts the dominant eigenvalues of a block-encoded Hermitian operator σ in time polylogarithmic in its dimension [8, 9]. Therefore, combining these two in the right way would seem to hand us a quantum algorithm for SPOD at polylogarithmic cost. This raises two questions.

*Mathematical Institute, University of Oxford, United Kingdom.

†Quantum Software Lab, University of Edinburgh, United Kingdom.

‡Mechanical Engineering, National University of Singapore, Singapore.

§Centre for Quantum Technologies, National University of Singapore & Inveriant, Singapore.

First, do the two subroutines compose in the right way? QFTs and qPCA-algorithms are each efficient in isolation, but qPCA diagonalizes a *single* Hermitian operator, whereas the QFT leaves the frequency resolved data spread across a register, with SPOD solving a separate spectral problem at each frequency. The construction thus hinges on whether there is a single operator that carries all of SPOD and admits an efficient block-encoding.

Second, does any resulting speedup survive dequantization? Quantum-inspired sample-and-query (SQ) techniques dequantize qPCA-style algorithms by exploiting low rank [10, 11, 12]. However, SQ access is not closed under the Fourier transformation, since each coefficient $\hat{Q}_{mlk}$ mixes all N_f time samples. Whether, and at what cost, SQ access propagates through the DFT is therefore not immediate.

Our contribution is to answer both questions. In Section 2.1 we propose a new quantum algorithm, qSPOD, which combines a QFT with a qPCA subroutine to find all dominant triples $(\lambda_{k,r}, k, |\psi_{k,r}\rangle)$ using a number of queries to a state-preparation unitary that is polylogarithmic in M , L , and N_f , yielding an exponential separation from the classical cost of Eq. (1). In Section 2.2 we show this advantage is only partial: a new classical algorithm, dSPOD, assuming efficient SQ access to Q matches the polylogarithmic dimension-scaling, leaving us with only (large) polynomial overheads. In Section 3 we compare all three—classical SPOD, qSPOD, and dSPOD—and outline directions for future work.

2 Results

To formalize our results we first define the problem and state our assumptions on the data and the access model. Without loss of generality we assume that the input data is normalized, i.e. $\|Q\|_F^2 = 1$, which is preserved by the unitary DFT. Therefore, the SPOD energies satisfy $\sum_{k,r} \lambda_{k,r} = 1$, and $\lambda_{k,r} \in (0, 1)$ is the fraction of total energy carried by mode $\psi_{k,r}$.

Problem 1 (SPOD). *Let $Q \in \mathbb{R}^{M \times L \times N_f}$ with $\|Q\|_F^2 = 1$, and $\epsilon, \tau \in (0, 1)$ with $\epsilon \leq \tau$. Given access to Q , output every triple $(\lambda_{k,r}, k, |\psi_{k,r}\rangle)$ with $\lambda_{k,r} \geq \tau$, reporting each energy $\lambda_{k,r}$ to additive error at most ϵ .*

While Problem 1 returns the eigenstates $|\psi_{k,r}\rangle$, it certifies only the energies $\lambda_{k,r}$. A fidelity guarantee on the eigenstates is a strictly stronger requirement, deferred to the full version.

Assumption 1 (Energy concentration). *Let τ be the threshold parameter of Problem 1. We assume that a threshold $\tau \geq 1/\text{polylog}(M, N_f)$ captures a constant fraction of the total energy, i.e.*

$$\sum_{k,r} \lambda_{k,r} \mathbf{1}[\lambda_{k,r} \geq \tau] \geq 1 - \delta,$$

for a constant $\delta \in [0, 1)$. As a consequence of $\sum_{k,r} \lambda_{k,r} = 1$, the number of triples above the threshold is bounded by $1/\tau = O(\text{polylog}(M, N_f))$.

Assumption 2 (Access). *For qSPOD, we assume access to a unitary U_Q with $|Q\rangle := U_Q |0\rangle = \sum_{m,l,t} Q_{mlt} |m\rangle |l\rangle |t\rangle$, together with its inverse $U_Q^\dagger$. For dSPOD, we assume $O(T)$ -time sample-and-query (SQ) access to Q [11].*

Assumption 2 can be satisfied by the usual QRAM model [13, 14]), but is in fact slightly weaker. qSPOD requires only a state-preparation unitary (and its inverse) for $|Q\rangle$, not the full block-encoding U_Q that QRAM provides. In particular, qSPOD also admits quantum native inputs opening new interesting avenues. Assumption 1 is the physical regime of interest, i.e. turbulent flows where a few modes capture the dominant dynamics.

2.1 Quantum SPOD

We now answer the first question by naming the operator σ and constructing a block-encoding U_σ of it. Starting from $|Q\rangle = \sum_{m,l,t} Q_{mlt} |m\rangle |l\rangle |t\rangle$, we apply a quantum Fourier transform to the last register to obtain $|\hat{Q}\rangle = \sum_{m,l,k} \hat{Q}_{mlk} |m\rangle |l\rangle |k\rangle$, where $\hat{Q}_{mlk} = (\hat{Q}_k)_{ml}$ with $\hat{Q}_k$ as defined before; a single QFT of cost $O(\text{polylog } N_f)$ thus realizes all $M \times L$ block transforms at once. Discarding the block register and dephasing the frequency register leaves the block-diagonal density matrix

$$\sigma = \bigoplus_{k=0}^{N_f-1} C_k,$$

the single operator that carries SPOD. Because σ is block diagonal, its spectrum is the disjoint union $\bigcup_k \{\lambda_{k,r}\}_{r=1}^R$ and its eigenvectors $|k\rangle \otimes |\psi_{k,r}\rangle$ are self-labeled by frequency. Hence, a single qPCA routine on σ can recover the dominant triples $(\lambda_{k,r}, k, |\psi_{k,r}\rangle)$ across all N_f frequencies at once. Since $|\hat{Q}\rangle$ (with the dephasing ancilla) purifies σ , Lemma 45 of [15] yields a $(1, 2n_K + n_L + n_M, 0)$ -block-encoding U_σ of σ

using $O(1)$ queries to $U_Q, U_Q^\dagger$ and $O(\text{polylog } N_f)$ additional gates, where n_K, n_L, n_M count the qubits of the frequency, block, and spatial registers.

Access to the block-encoding U_σ and copies of the purification $|\hat{Q}\rangle$ lets us perform qPCA-style routines on σ . The full manuscript introduces a dedicated qPCA subroutine that uses windowed QPE [16] to guarantee high-fidelity eigenvectors and resolve degeneracies, here we use the simpler singular-value estimation (SVE) of Theorem 27 of [8], which suffices for Problem 1. Since U_σ is exact the SVE’s input-accuracy requirement is trivially met, and measuring the SVE register tags each eigenvector with its energy $\lambda_{k,r}$. Together with the self-labeling frequency register this collapses the purification onto the triple $(\lambda_{k,r}, k, |\psi_{k,r}\rangle)$ with probability $\lambda_{k,r}$. Since at most $1/\tau$ modes exceed τ and each is sampled with probability at least τ , a coupon-collector argument recovers all required triples after $\tilde{O}(1/\tau)$ repetitions.

Result 1 (qSPOD). *There exists a quantum algorithm, qSPOD, that solves Problem 1 with high probability, using $\tilde{O}(\frac{1}{\epsilon\tau})$ queries to U_Q and $U_Q^\dagger$, where the notation $\tilde{O}$ hides factors polylogarithmic in M, L , and N_f .*

Under Assumption 1 we have $\tau \geq 1/\text{polylog}(M, N_f)$, so the cost of Result 1 reduces to $\tilde{O}(1/\epsilon)$. Moreover the recovered modes are then promised to carry a $\geq 1 - \delta$ fraction of the total energy. This is an exponential improvement in M, L , and N_f over Eq. (1). Notably, if classical descriptions of the dominant eigenmodes $\psi_{k,r}$ are required, one can combine eigenstate filtering [17] with state tomography [18], which reintroduces a $\text{poly}(M)$ cost per mode and thus forfeits the exponential advantage in M .

2.2 Dequantized SPOD

We now address the second question. $O(T)$ -time query access to the entries Q_{mlt} yields $O(N_f T)$ -time query access to the coefficients $(\hat{Q}_k)_{ml} = \hat{Q}_{mlk}$ by summing the N_f time entries the Fourier transform mixes. Any algorithm built on this dense access thus inherits a $\text{poly}(N_f)$ factor. If, however, each of the $M \times L$ time-series Q_{ml*} is (approximately) s -sparse in the Fourier basis, sparse FFT recovers the series’ spectrum, and hence query access to $\hat{Q}_{ml*}$ in time $\tilde{O}(sT)$ [19]. To our knowledge, this is the first time that this technique is used in the context of dequantized algorithms. Assuming $s \in O(\text{polylog}(N_f))$ removes the N_f -dependence and is of similar spirit as Assumption 1 in a sense that both assume that coherent flows are sparse in a suitable basis. To obtain sampling access to $\hat{Q}_k$ —i.e. sampling $(m, l) \sim |(\hat{Q}_k)_{ml}|^2 / \|\hat{Q}_k\|_F^2$, which, by discarding l , marginalizes to the row distribution $m \sim \|(\hat{Q}_k)_{m*}\|_2^2 / \|\hat{Q}_k\|_F^2$ —we use rejection sampling with the Parseval envelope $|(\hat{Q}_k)_{ml}|^2 \leq \|Q_{ml*}\|_2^2$. This incurs no overhead in M or L , only an acceptance overhead of $1/\|\hat{Q}_k\|_F^2 = 1/\text{tr}(C_k)$, which is at most $1/\tau$ for every frequency bin k carrying a mode $\lambda_{k,r} \geq \tau$. Combining these observations with the results from [12, 11] yields the following result.

Result 2 (dSPOD). *Given $O(T)$ -time SQ access to Q with s -sparse time-series Q_{ml*} , there exists a classical algorithm, dSPOD, that solves Problem 1 with high probability, in time $\tilde{O}(sT \text{poly}(\frac{1}{\epsilon}, \frac{1}{\tau}))$, where the notation $\tilde{O}$ hides factors polylogarithmic in M, L , and N_f .*

3 Conclusion

	Runtime	Access model	Output
SPOD	$\tilde{O}(N_f M L^2)$	Classical (RAM) on Q	Full spectrum $\{(\lambda_{k,r}, k, \psi_{k,r})\}$
qSPOD	$\tilde{O}(\frac{1}{\epsilon\tau})$ (Result 1)	Unitaries $U_Q, U_Q^\dagger$	Top spectrum $\{(\lambda_{k,r}, k, \psi_{k,r}\rangle)\}$
dSPOD	$\tilde{O}(s \text{poly}(\frac{1}{\epsilon}, \frac{1}{\tau}))$ (Result 2)	SQ(Q)	Top spectrum $\{(\lambda_{k,r}, k, \text{SQ}(\psi_{k,r}))\}$

Table 1: Comparison of classical SPOD, quantum qSPOD, and dequantized dSPOD solving Problem 1 under Assumption 1. Top spectrum is with respect to the threshold parameter τ , and energies are returned to precision ϵ .

The above table summarizes the three algorithms. qSPOD attains a runtime polylogarithmic in M, L , and N_f under state-preparation access, an exponential improvement over the classical $\tilde{O}(M N_f L^2)$, by representing all frequency-wise spectral problems as a single block-diagonal operator. This exponential separation, however, does not survive dequantization. dSPOD, which assumes sample-and-query access, recovers the same dimension-scaling classically, at a polynomial overhead whose exponents are large enough that qSPOD may remain perfarble in practice. Natural next steps are a windowed-QPE qPCA subroutine with certified eigenvector fidelity and multiplicity resolution, and an exploration of quantum processes as native input candidates for qSPOD.

References

- [1] John Leask Lumley. The structure of inhomogeneous turbulent flows. *Atmospheric turbulence and radio wave propagation*, pages 166–178, 1967.
- [2] John L. Lumley. *Stochastic Tools in Turbulence*. Academic Press, New York, 1970.
- [3] Aaron Towne, Oliver T Schmidt, and Tim Colonius. Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *Journal of Fluid Mechanics*, 847:821–867, 2018.
- [4] Oliver T Schmidt and Tim Colonius. Guide to spectral proper orthogonal decomposition. *Aiaa journal*, 58(3):1023–1033, 2020.
- [5] Oliver T Schmidt, Gianmarco Mengaldo, Gianpaolo Balsamo, and Nils P Wedi. Spectral empirical orthogonal function analysis of weather and climate data. *Monthly Weather Review*, 147(8):2979–2995, 2019.
- [6] Oliver T Schmidt, Aaron Towne, Tim Colonius, André VG Cavalieri, Peter Jordan, and Guillaume A Brès. Wavepackets and trapped acoustic modes in a turbulent jet: coherent structure eduction and global stability. *Journal of Fluid Mechanics*, 825:1153–1181, 2017.
- [7] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010.
- [8] Shantanav Chakraborty, András Gilyén, and Stacey Jeffery. The power of block-encoded matrix powers: improved regression techniques via faster Hamiltonian simulation. *arXiv preprint arXiv:1804.01973*, 2018.
- [9] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nature physics*, 10(9):631–633, 2014.
- [10] Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*, pages 217–228, 2019.
- [11] Ewin Tang. Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions. *Physical Review Letters*, 127(6):060503, 2021.
- [12] Nai-Hui Chia, András Pal Gilyén, Tongyang Li, Han-Hsuan Lin, Ewin Tang, and Chunhao Wang. Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning. *Journal of the ACM*, 69(5):1–72, 2022.
- [13] Connor T Hann, Gideon Lee, SM Girvin, and Liang Jiang. Resilience of quantum random access memory to generic noise. *Prx Quantum*, 2(2):020311, 2021.
- [14] Alexander M Dalzell, András Gilyén, Connor T Hann, Sam McArdle, Grant Salton, Quynh T Nguyen, Aleksander Kubica, and Fernando GSL Brandão. A distillation-teleportation protocol for fault-tolerant qram. *arXiv preprint arXiv:2505.20265*, 2025.
- [15] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. *arXiv preprint arXiv:1806.01838*, 2018.
- [16] Dominic W Berry, Yuan Su, Casper Gyurik, Robbie King, Joao Basso, Alexander Del Toro Barba, Abhishek Rajput, Nathan Wiebe, Vedran Dunjko, and Ryan Babbush. Analyzing prospects for quantum advantage in topological data analysis. *PRX Quantum*, 5(1):010319, 2024.
- [17] Lin Lin and Yu Tong. Optimal polynomial based quantum eigenstate filtering with application to solving quantum linear systems. *Quantum*, 4:361, 2020.
- [18] Joran van Apeldoorn, Arjan Cornelissen, András Gilyén, and Giacomo Nannicini. Quantum tomography using state-preparation unitaries. In *Proceedings of the 2023 annual ACM-SIAM symposium on discrete algorithms (SODA)*, pages 1265–1318. SIAM, 2023.
- [19] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric Price. Nearly optimal sparse fourier transform. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 563–578, 2012.