

An efficient state preparation for the Pauli correlation encoding

Lucas Borges,^{1,2,*} Marco Sciorilli,^{1,3,*} and Leandro Aolita¹

¹*Quantum Research Center, Technology Innovation Institute, Abu Dhabi, UAE*

²*Federal University of Rio de Janeiro, Caixa Postal 652, Rio de Janeiro, RJ 21941-972, Brazil*

³*Laboratoire d'Informatique de Paris 6, CNRS,
Sorbonne Université, 4 Place Jussieu, 75005 Paris, France*

Pauli correlation encoding (PCE) compresses many classical variables into expectation values of Pauli strings on fewer qubits. This compression has been used as a variational representation for combinatorial optimization, but a complementary primitive is needed to make the encoding bidirectional: given classical data or a classical solution, one would like to upload it directly into the compressed Pauli-correlation register. We introduce an explicit state-preparation unitary for this task. For the family of k -body Z correlators, we prove that the induced correlation map has full-rank Jacobian at the origin and admits a certified local inverse. In particular, for $m = \binom{n}{k}$ encoded coordinates, every target vector in an explicit image cube of radius $\rho(r) \in \Omega(1/m)$ can be prepared from a parameter cube with $r < \log(2)/m$. We then use this upload primitive in two ways. First, we warm-start PCE optimization by injecting classical bit strings into the signs of Pauli correlators before a second variational optimization round. On weighted MaxCut instances with 400 variables encoded on 15 qubits, this improves the solutions obtained after the first round and suppresses the barren-plateau behavior observed for standard random initialization. Second, we demonstrate classical data upload on Quantinuum's H2-2 trapped-ion processor by encoding and retrieving images with 494 and 1476 pixels using only two-body ZZ correlations on 32 and 56 qubits, respectively.

Main idea. Pauli correlation encoding (PCE) [1] represents classical variables through signs or values of Pauli expectation values. Given a set of Pauli strings $\Pi = \{\Pi_i\}_{i=1}^N$, a binary assignment can be decoded as

$$x_i = \text{sgn}(\langle \Pi_i \rangle), \quad i \in [N], \quad (1)$$

where $\langle \Pi_i \rangle = \langle \psi | \Pi_i | \psi \rangle$. This allows N classical variables to be stored in the correlations of an n -qubit state, with N polynomially larger than n when the Pauli set is chosen from k -body correlators. Previous uses of PCE focused on variationally finding a quantum state whose measured signs give a good solution to an optimization problem [1, 2]. Here we study the reverse direction: given a classical vector of target correlations, can we prepare a quantum state whose Pauli expectations realize it?

For the standard PCE family of k -body Z correlators, we give an explicit product of Pauli rotations that maps the initial state $|+\rangle^{\otimes n}$ to a state whose correlators can be prescribed independently in a certified neighborhood of the origin. This primitive has two immediate uses. In optimization, it provides a warm-start layer [3]: a classical bit string can be injected into the signs of the PCE correlators before a new variational round. In quantum machine learning, it provides a data-loading layer [4]: a real-valued feature vector, such as a normalized image, can be uploaded into a compressed quantum register as a vector of Pauli correlations rather than as amplitudes.

The state-preparation unitary. Let $\Omega_k = \{I \subset [n] : |I| = k\}$ and $m = |\Omega_k| = \binom{n}{k}$. For each $I \in \Omega_k$, define the PCE observable

$$\Pi_I = \prod_{a \in I} Z_a. \quad (2)$$

The upload circuit is

$$U(\theta) = \prod_{I \in \Omega_k} \exp\left(-\frac{i}{2}\theta_I Q_I\right), \quad \theta \in \mathbb{R}^m, \quad (3)$$

with $Q_I = iX_{\pi(I)}\Pi_I = Y_{\pi(I)}\prod_{a \in I \setminus \{\pi(I)\}} Z_a$, where $\pi(I) \in I$ is the pivot of the set I .

The induced PCE correlation map is

$$f_I(\theta) = \langle + |^{\otimes n} U(\theta)^\dagger \Pi_I U(\theta) | + \rangle^{\otimes n}. \quad (4)$$

* Both authors contributed equally to this work




FIG. 1: **Performance and Barren Plateau.** **Picture a.** Performance of the warm-start PCE on 10 different weighted random MaxCut instances with 400 variables and integer edge weights between -10 and 10, encoded on 15 qubits using cubic compression. We report the improvement in approximation ratio after a second round of optimization as a function of the approximation ratio at the end of the first round. For comparison, we show the same metric for another competitive iterative method, CirCut, which is based on the Burer-Monteiro algorithm. For easier visualization, we grouped the data points into intervals of width 0.1 and displayed their distributions as violin plots. On the tested instances, the warm-start PCE appears to outperform Burer-Monteiro, with a gap that increases with the starting approximation ratio. **Picture b.** Average variance of the loss function per parameter computed over 100 random initializations, over the whole gradient and 100 random weighted MaxCut problems, as a function of the number of qubits used to encode the problem. The size of the MaxCut problem grows quadratically with the number of qubits, as we used all the expectation values of the two-body (ZZ) Pauli correlators. In log scale, we compare the decay of the standard PCE against the warm-started one presented in this work. We observe that, while the standard PCE shows the expected exponential decay, the variance of the warm-started variant quickly plateaus.

The generators are chosen so that, by construction, each Q_I anticommutes with its associated observable Π_I , and the Jacobian $Df(0) = -\mathbb{1}_m$ is full-rank, guaranteeing that infinitesimally, every encoded correlator is independently controlled by its corresponding parameter [5].

We prove that any target feature vector $y \in [-1, 1]^m$, such that $\|y\|_\infty \leq \rho(r)$, with $\rho(r) \in \Omega(1/m)$, can be uploaded into the PCE register. Concretely, there is a unique $\theta \in [-1, 1]^m$ with $\|\theta\|_\infty \leq r$, $r \in \Omega(1/m)$, such that

$$\langle + |^{\otimes n} U(\theta)^\dagger \Pi_I U(\theta) | + \rangle^{\otimes n} = y_I, \quad \forall I \in \Omega_k. \quad (5)$$

This is not merely an existence result. Inside the certified cube, a simple Euler method is guaranteed to converge to the upload parameters. The evaluations of f and Df are also classically approximable to any desired precision in this regime, for instance by truncating the Pauli-rotation Taylor expansion or by using standard circuit-approximation methods for products of small Pauli rotations. Thus the upload layer is an efficiently computable classical preprocessing step.

For warm-start optimization, one uploads a scaled bit string $y_I = \lambda b_I$, with $b_I \in \pm 1$ and $0 < \lambda \leq \rho(r)$, so that the signs of the measured correlators reproduce the desired assignment. For data loading, one similarly rescales a classical feature vector $x \in [-1, 1]^m$ and uploads $y_I = \lambda x_I$.

PCE as an iterative method. Many successful algorithms for combinatorial optimization operate iteratively. Starting from a candidate solution, they introduce a controlled perturbation and restart the optimization in its neighborhood [6–9]. The procedure is intended to escape the local minimum reached during the preceding run while retaining information contained in the current solution.

We apply the same iterative mechanic to quantum optimization based on the Pauli correlation encoding. As introduced in Ref. [1], we use the signs of a selected set of Pauli correlators to encode a binary assignment for a QUBO problem. Using the state-preparation procedure described above, we initialize these correlators with the signs associated with a solution obtained in a preceding optimization round. A perturbation is then introduced by randomly sampling their amplitudes within a prescribed range, thereby preserving the encoded binary solution while displacing the initial quantum state within its neighborhood. We then append a parametrized circuit to the prepared state, initialized so that it acts as the identity. At the beginning of the second optimization, the measured correlations therefore coincide with those of the uploaded solution. A classical optimizer is then used to find the optimal parameters of the parametrized circuit, in the standard VQA fashion, minimizing once again the loss function defined as in Ref. [1]. The resulting solution can in turn be re-encoded and used to initialize a further optimization round.

In Fig. 1(a) we numerically evaluate the warm-start protocol on weighted random MaxCut instances containing 400 variables, with integer edge weights between -10 and 10. In the first optimization round, a standard VQA using PCE with




FIG. 2: **Classical data upload on a quantum processing unit.** Images uploaded to the Quantinuum H2-2 trapped-ion QPU using only two-body ZZ Pauli correlators. **Picture a.** The image was encoded using 494 expectation values over 32 qubits, with pixel intensities discretized into 7 color intervals between $-1/4$ and $1/4$. From left to right, the panels show the original image, the image retrieved from the hardware using approximately 6×10^3 shots, and the image obtained by augmenting the hardware data with an additional 6×10^3 shots from the Quantinuum H2-2 noisy emulator. The emulator-assisted reconstruction shows no pixel errors. **Picture b.** Magritte’s *The Son of Man* was encoded using 1476 expectation values over 56 qubits, with pixel colors discretized into 8 color intervals between $-1/6$ and $1/6$. From left to right, the panels show the original image, the pixelated image used as the uploaded data, and the image retrieved from the hardware using approximately 8×10^3 shots. At this system size, noisy emulation of the H2-2 device was no longer feasible. The retrieved image has a pixel error of approximately 21%.

cubic compression on 15 qubits is used to obtain an initial solution. The perturbed solution is injected into a quantum state using the same encoding setup, and a second round of optimization on the parametrized ansatz is then initialized from this state. As a classical reference, we consider CirCut, an iterative method based on the Burer–Monteiro rank-2 relaxation algorithm.

Across the tested instances, warm-started PCE yields a greater average improvement after the second optimization than the classical baseline. The difference is most pronounced for initial solutions that are already close to the best-known value. Such solutions are associated with deeper local minima and are therefore difficult to improve through an unstructured restart.

These strong results also suggest a hybrid quantum–classical strategy in which a classical solver identifies a high-quality solution, while warm-started PCE performs the final “last-mile” optimization required to escape the deep local minima and refine it further.

Furthermore, the small rotation angles used in the state-preparation circuit, together with the identity initialization of the variational ansatz, keep the circuit close to the initial $|+\rangle^{\otimes n}$ state at the start of the optimization. The resulting unitary ensemble remains far from a unitary t -design, and the concentration-of-measure mechanism associated with barren plateaus is therefore not expected to arise at initialization. In Fig. 1(b), we report the gradient variance for comparison. Whereas standard PCE exhibits the expected exponential suppression of the gradient variance, the warm-started initialization reaches a finite plateau that remains approximately constant with increasing system size.

Data upload on trapped ions. A further application of the state-preparation method is the upload of classical data for quantum processing. We tested the viability of it on current quantum hardware using two images containing 494 and 1476 pixels, encoded in the expectation values of ZZ correlators on 32 and 56 qubits, respectively. The allowed range of each correlator was divided into eight intervals, which were mapped linearly to the image color values. After preparing the state, we measured the correlators and decoded the image. With sufficient measurement statistics, the smaller image was recovered exactly. For the larger image, hardware noise produced a pixel error of approximately 21%, although all incorrect pixels were assigned to an adjacent color interval. All experiments were performed on the Quantinuum H2-2 processor. The initial $|+\rangle^{\otimes n}$ state was prepared using Hadamard gates, followed by one two-qubit R_{ZZ} rotation for each encoded pixel. The results are shown in Fig. 2. With the base-8 encoding, the largest experiment corresponds to the upload and recovery of 4,428 bits of classical data, which, to the best of our knowledge, is the largest such demonstration reported to date.

Discussion. We introduced a state-preparation primitive for Pauli correlation encoding. The mathematical result is a certified and algorithmic local inverse for the PCE correlation map: for $m = \binom{n}{k}$ encoded k -body Z correlators, every target vector in a cube of radius $\rho(r) \in \Omega(1/m)$ can be realized by an explicit product of Pauli rotations on n qubits. This guarantee is conservative, and the observed sparsity of the nonlinear terms suggests that the optimal upload radius may be substantially larger.

Our numerical warm-start experiments indicate that this structure can improve optimization after an initial PCE round and can mitigate the gradient-variance decay associated with random initialization. Our trapped-ion experiments show that the same preparation idea can be deployed on present hardware to upload and retrieve nontrivial classical images from Pauli correlators on up to 56 qubits. Together, these results suggest that correlation-based state preparation may be a useful near-term primitive for both combinatorial optimization and data-driven quantum algorithms.

-
- [1] M. Sciorilli, L. Borges, T. L. Patti, D. García-Martín, G. Camilo, A. Anandkumar, and L. Aolita, Towards large-scale quantum optimization solvers with few qubits, *Nature Communications* **16**, 476 (2025).
 - [2] M. Sciorilli, G. Camilo, T. O. Maciel, A. Canabarro, L. Borges, and L. Aolita, A competitive nisq and qubit-efficient solver for the labs problem (2026), arXiv:2506.17391 [quant-ph].
 - [3] D. J. Egger, J. Mareček, and S. Woerner, Warm-starting quantum optimization, *Quantum* **5**, 479 (2021).
 - [4] O. Zang, G. Barraué, and T. Quertier, Benchmarking data encoding methods in quantum machine learning (2025), arXiv:2505.14295 [quant-ph].
 - [5] V. I. Arnol'd, *Mathematical Methods of Classical Mechanics*, 2nd ed., Graduate Texts in Mathematics, Vol. 60 (Springer-Verlag, New York, NY, 1989).
 - [6] P. Festa, P. Pardalos, M. Resende, and C. R. and, Randomized heuristics for the max-cut problem, *Optimization Methods and Software* **17**, 1033 (2002), <https://doi.org/10.1080/1055678021000090033>.
 - [7] E. Arráiz and O. Olivo, Competitive simulated annealing and tabu search algorithms for the max-cut problem, in *Proceedings of the 11th Annual Conference on Genetic and Evolutionary Computation, GECCO '09* (Association for Computing Machinery, New York, NY, USA, 2009) p. 1797–1798.
 - [8] G. A. Kochenberger, J.-K. Hao, Z. Lü, H. Wang, and F. Glover, Solving large scale max cut problems via tabu search, *Journal of Heuristics* **19**, 565 (2013).
 - [9] U. Benlic and J.-K. Hao, Breakout local search for the max-cutproblem, *Eng. Appl. Artif. Intell.* **26**, 1162–1173 (2013).