Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Quantum approximate optimization of population-level acetabular screw parameters on gate-based hardware

Günbay Noyan Dirlik[1]

ndirlik@gmail.com

Alanya Alaaddin Keykubat Üniversitesi

Article

Keywords: Acetabular fractures, retrograde screw fixation, quantum computing, combinatorial optimization, QAOA, implant selection

Posted Date: September 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10821375/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Quantum optimization has been proposed for biomedical decision problems, but its behaviour on current hardware remains uncertain. We formulated a simplified population-level acetabular screw parameter-selection task from published CT morphometry as a quadratic unconstrained binary optimization (QUBO) problem and evaluated the Quantum Approximate Optimization Algorithm (QAOA). A reduced five-candidate encoding was run at depths p = 1–4 on two IBM Heron r2 processors with repeated 1,000-shot experiments, while a complete 14-qubit encoding was assessed in noiseless simulation. The best hardware run sampled the model optimum with probability 0.409 ± 0.008, about thirteen times the 1/32 random baseline, although classical methods solved the same instance exactly in milliseconds. Across eight population–sex profiles, hardware success probabilities ranged from 0.223 to 0.334. The complete encoding sampled the optimum with probability 0.019, about three times uniform sampling over admissible configurations, but showed marked seed dependence. These findings show that one-hot cardinality constraints can make variational optimization difficult even when the underlying clinical objective is separable and classically simple, and define current limitations for translation to patient-specific planning.

Introduction

Percutaneous retrograde intramedullary screw fixation is used for selected anterior-column acetabular fractures as a less invasive alternative to open approaches, with potential reductions in blood loss and wound morbidity [1–4]. Its feasibility depends on an osseous corridor whose dimensions vary between individuals, sexes, and populations [5–9].

Reported minimum corridor diameters vary substantially across populations, from 7.35 mm in Arab males to 12.33 mm in Turkish males, with corresponding differences in recommendations for 6.5-mm screws [5–9]. Our comparative study found a larger minimum diameter in Turkish than Russian males (12.33 vs 10.69 mm, p = 0.019) [5]. Entry-point studies also describe a crescent-shaped admissible region rather than a single reproducible point [10], and minimum diameter alone cannot represent a curved three-dimensional trajectory [11].

Clinically, entry point, trajectory, length, and diameter are anatomically coupled. This population-level approximation instead treats corridor diameter and maximum length as fixed parameters, making the objective separable. It is not a patient-specific planning system, and this simplification is central to interpretation.

QAOA is a variational algorithm for combinatorial optimization [12, 13]. Current processors operate in the NISQ regime, where circuit depth and accumulated error limit performance [14]. Quantum computing in orthopaedic and surgical literature remains dominated by perspectives [15–18]. A recent acetabular framework combined quantum and biomechanical concepts but did not report hardware execution [24]. To our knowledge, no orthopaedic parameter-selection QUBO has previously been executed on gate-based quantum hardware.

The study therefore asks a limited question: can a simplified population-level orthopaedic parameter-selection problem be formulated in a quantum-compatible representation, executed on current hardware, and compared transparently with classical ground truth?

Results

Classical optimum and shortlist

All 162 Turkish male configurations satisfied the hard constraints. The selected 6.5-mm diameter was consistent with the recommendation in the source morphometric study, providing a face-validity check rather than independent clinical validation [5].

Table 2. Classical optimum and five lowest-cost configurations used in the 5-qubit QAOA benchmark.

Rank	Entry	Sagittal	Horizontal	Length	Diameter	Cost	Normalized
0	26.7 mm	54°	33°	125 mm	6.5 mm	0.479	0.000
1	26.7 mm	54°	33°	120 mm	6.5 mm	0.559	0.330
2	26.7 mm	54°	33°	115 mm	6.5 mm	0.639	0.659
3	26.7 mm	54°	33°	125 mm	4.5 mm	0.641	0.670
4	26.7 mm	54°	33°	120 mm	4.5 mm	0.721	1.000

Entry point and both angles remained at their population means throughout the shortlist. Length was the main discriminator, followed by diameter.

Penalty selection

Table 3. Penalty-coefficient check for the 5-qubit QUBO.

A	All-zero	Cand. 0	Cand. 1	Cand. 2	Cand. 3	Cand. 4	All invalid states above valid range?
0.50	0.500	0.000	0.330	0.659	0.670	1.000	No
1.00	1.000	0.000	0.330	0.659	0.670	1.000	No
1.50	1.500	0.000	0.330	0.659	0.670	1.000	Yes
2.00	2.000	0.000	0.330	0.659	0.670	1.000	Yes
3.00	3.000	0.000	0.330	0.659	0.670	1.000	Yes

At A ≤ 1, the all-zero state lies within the valid energy range, although candidate 0 remains the global minimum. Any A > 1 therefore places every cardinality-violating state above all valid one-hot states; within the coefficients tested, this criterion was first met at A = 1.5, and A = 2.0 was adopted to provide margin.

QAOA performance versus circuit depth

Table 4. Probability of sampling the optimum by QAOA depth (four repeats × 1000 shots; mean ± SE).

p	Transpiled depth	CZ gates	Simulator	ibm_fez	ibm_marrakesh
1	106–111	35	0.311	0.327 ± 0.008	0.298 ± 0.010
2	212–225	70	0.469	0.409 ± 0.008	0.311 ± 0.011
3	323–330	114	0.453	0.293 ± 0.005	0.255 ± 0.007
4	415–435	152	0.744	0.307 ± 0.004	0.382 ± 0.003

Within-job variability was low across the four repeats (inter-repeat standard deviations 0.009-0.015). On ibm_fez, hardware retention relative to the simulator was 105%, 87%, 65%, and 41% at p = 1-4, decreasing monotonically as CZ count increased from 35 to 152. At p = 1, the hardware and simulator estimates were within shot-level uncertainty; at p = 4, hardware retained approximately 41% of simulator probability.

The depth ordering was not robust across devices. On ibm_fez, p = 4 performed below p = 1, whereas on ibm_marrakesh p = 4 produced the highest hardware value. Because the devices were tested at different times, device effects cannot be separated from temporal variation. No optimal depth or device superiority is claimed.

Simulator expected energy decreased with p (1.232, 0.766, 0.649, 0.583), whereas optimal-state probability was nonmonotonic (0.311, 0.469, 0.453, 0.744). These objectives are not equivalent. Because the high p = 4 simulator probability directly entered the noise estimate, that comparison should be interpreted cautiously.

Simplified noise budget

The estimated readout-survival factor was 0.959, insufficient to explain depth-dependent loss. The observed/model ratio fell from 1.21 at p = 1 to 0.66 at p = 4, showing that the independent-error approximation did not capture the experimental behaviour. The data cannot distinguish heterogeneous gate errors, idle errors, crosstalk, coherent effects, calibration drift, or other depth-dependent mechanisms.

Behaviour across population profiles

The same rule set was applied to Turkish, Russian, Chinese [7], and German [8] male and female profiles with complete diameter, length, and angular data. Arab [6] and Indian [9] cohorts were not imputed because the required angular parameters were unavailable. This analysis was a consistency check rather than a new clinical finding.

Table 5. Classical selection and hardware performance by population profile at p = 2 on ibm_marrakesh (four repeats × 1000 shots; mean ± SE).

Profile	Min corridor (mm)	Valid	Diameter (mm)	Simulator	Hardware	Conditional
Turkish M	12.33	162/162	6.5	0.469	0.331 ± 0.007	0.573
Turkish F	11.52	162/162	6.5	0.466	0.334 ± 0.018	0.581
Russian M	10.69	162/162	6.5	0.443	0.312 ± 0.007	0.547
Russian F	11.82	162/162	6.5	0.465	0.328 ± 0.007	0.572
Chinese M	8.70	162/162	6.5	0.435	0.303 ± 0.006	0.540
Chinese F	7.60	81/162	4.5	0.301	0.223 ± 0.007	0.392
German M	9.50	162/162	6.5	0.449	0.303 ± 0.011	0.536
German F	7.50	81/162	4.5	0.303	0.232 ± 0.007	0.397

In the Chinese and German female profiles, a 6.5-mm screw left less than the prespecified 1-mm cortical margin; all 6.5-mm configurations were therefore excluded and the model selected 4.5 mm. This reproduces the direction of population-specific recommendations using one unchanged rule set, but follows directly from the imposed margin threshold.

All eight circuits used the same physical qubits, depth 212, and 70 CZ gates. Hardware retained 67%-76% of simulator probability without a corridor-width trend. Narrow-corridor profiles had lower simulator probabilities because min-max normalization compressed three near-degenerate length options relative to the A = 2 penalty. Conditional success was 39%-40% in these profiles versus 54%-58% in wider corridors (20% baseline). The Turkish male values in the depth and profile experiments were acquired in different sessions, consistent with the session dependence noted below.

Non-reproduction of an earlier session

An earlier session produced lower values (0.201-0.223 for wide corridors) and an anomalous Chinese female result of 0.010. A later controlled run gave 0.223 ± 0.007 for that profile and a Turkish male value consistent with two independent measurements. The earlier session was therefore not used as the principal analysis. This non-reproduction illustrates the need for internal controls and circuit-level replication.

Measurement-error mitigation

Measurement mitigation produced no detectable improvement (raw 0.309 ± 0.005; mitigated 0.308 ± 0.006). Readout error was therefore unlikely to be the dominant limitation, although no single mechanism could be isolated. Device and time effects remained confounded in the two-device depth sweep.

Robustness of cost weights

Each cost weight was varied individually by ±50% and ±100%, and 2000 random weight vectors were generated by independently scaling each weight from 0.5× to 2×.

The selected configuration was unchanged under all individual and random weight perturbations. This robustness was structural rather than empirical validation: mean entry point and angles had zero deviation, while the longest and largest admissible screw was monotonically rewarded. Anatomically coupled models would provide a more meaningful test of weighting.

Classical baselines

Table 6. Classical methods and QAOA on the primary 162-configuration instance.

Method	Reported compute time	Result
Tabu-style steepest descent	0.091 ms	Optimum in 100/100 starts
Exhaustive enumeration	0.146 ms	Exact optimum
Simulated annealing	4.6 ms	Optimum in 100/100 starts
QAOA, p = 2, ibm_fez	~3 s charged QPU time + queue	Optimal-state probability 0.409 ± 0.008
QAOA, p = 2, ibm_marrakesh	~3 s charged QPU time + queue	Optimal-state probability 0.311 ± 0.011

All classical methods located the optimum in under 5 ms. The QAOA time values are charged QPU time only and exclude queueing, statevector parameter optimization, and transpilation; Table 6 is therefore a workflow comparison rather than a complete time-to-solution benchmark.

Complete 14-qubit encoding

To remove classical pre-selection, the complete discretized instance was encoded with variable-wise one-hot groups: three qubits each for entry point, craniocaudal angle, mediolateral angle, and length, and two qubits for diameter. The resulting 14-qubit QUBO used a separate exactly-one penalty for each group. Exact enumeration confirmed the same global optimum as the classical model.

At p = 3 in noiseless simulation, the optimum was sampled with probability 0.019. Relative to uniform sampling over all 2^{14} bitstrings (1/16384 = 0.000061), this is approximately 300-fold. Because most bitstrings violate the one-hot constraints, the more informative baseline is uniform sampling among the 162 admissible configurations (1/162 = 0.0062), against which the result is approximately threefold. The objective decomposes into five independent lookups and is not combinatorially hard. Difficulty arose from five cardinality penalties around small candidate-cost differences. Across five seeds (15 restarts each), p = 3 success probabilities ranged from 0.005 to 0.027 (mean 0.016 ± 0.009). Because instability occurred without hardware noise, the 14-qubit circuit was not submitted to hardware. A nine-qubit

domain-wall attempt was excluded because its implemented energy representation was not exactly quadratic.
Thus, the 5-qubit benchmark is a stable but heavily reduced selection task, whereas the 14-qubit model represents the complete discretized instance without pre-selection. Its instability localizes the obstacle to one-hot penalty structure rather than qubit count, hardware noise, or the clinical objective itself.

Discussion

This proof-of-concept carried a simplified population-level orthopaedic parameter-selection model from classical morphometry through QUBO construction to gate-based hardware. The best hardware result sampled the model optimum with probability 0.409 ± 0.008, approximately thirteen times the 1/32 unconditional baseline, but still in fewer than half of shots. Classical enumeration and standard heuristics solved the same instance exactly and orders of magnitude faster. No quantum advantage was demonstrated or expected.

The main transferable finding concerns representation. Normalized candidate costs occupied a unit range, while cardinality penalties had to exceed that range. In the complete encoding, five penalties destabilized optimization despite a separable objective. This is representation-induced difficulty, not evidence that orthopaedic problems are intrinsically hard for QAOA. Binary encodings and feasible-subspace mixers warrant evaluation; Hamming-weight-preserving XY mixers have been executed on trapped-ion hardware for a comparable constrained selection problem [25].

The repeated hardware experiments also show that plausible single-session patterns may be misleading. An apparent depth trend and an anomalous population-profile result did not survive replication and internal bit-order checks. For clinically motivated NISQ studies, repeated shots alone are insufficient; repeated circuits, session controls, transparent reporting of superseded analyses, and comparison with noiseless ground truth are necessary.

Clinical relevance is prospective. Real entry point, trajectory, length, and diameter are coupled by three-dimensional anatomy, whereas this model treats them as separable and does not establish patient-specific planning or clinical safety. Its value is an auditable formulation and hardware benchmark. Automated CT corridor extraction may enable larger trajectory-coupled spaces [10, 19], but whether quantum optimization becomes competitive remains unknown.

Several limitations frame interpretation. The principal depth sweep used a single Turkish male profile and a five-candidate shortlist, and the two processors were tested at different times, preventing separation of device effects from temporal variation. The discretization was coarse and selected for qubit economy rather than anatomical resolution, and several clinically important constraints—including fracture reduction, fracture-line orientation, soft-tissue exit, and screw-head position—were not represented. Inputs were population means rather than patient-specific anatomy, while corridor diameter and maximum length were not coupled to the selected entry point and trajectory; the objective was therefore separable, the optimum dominant, and the cost weights not meaningfully discriminative. Absolute hardware probabilities were session-dependent, and min–max normalization depended on shortlist

composition. Two candidate-length constructions were used: rounded 5-mm values for the primary hardware instance and a profile-relative rule (L − 10, L − 5, L) for the cross-profile analysis; a single rule would have been cleaner but would have required repeating the eight-profile hardware session. Finally, the hard exclusions and cost weights reflected the judgement of a single surgeon rather than a consensus process or outcome data.

Population-level acetabular screw parameter selection can be expressed as a QUBO and executed on current gate-based quantum hardware. On the reduced 5-qubit instance, QAOA sampled the configuration consistent with the source morphometric recommendation with probability up to 0.409 ± 0.008, thirteen times the unconditional random baseline. On the complete 14-qubit instance in simulation, it reached approximately three times the baseline of uniform sampling among admissible configurations. Classical computation remained exact, reliable, and substantially faster.

No quantum advantage was observed. To our knowledge, this is the first reported execution of an orthopaedic parameter-selection QUBO on gate-based quantum hardware. The main methodological result is that cardinality-penalized one-hot encoding can impose a difficult variational landscape on an otherwise separable and computationally simple objective. Future work should prioritize patient-specific three-dimensional coupling and more efficient feasibility-preserving encodings rather than scaling the present model unchanged.

This proof-of-concept was intended to move quantum computing in orthopaedics from speculation toward a reproducible hardware experiment. It demonstrates neither clinical utility nor quantum advantage, but provides a transparent, testable account of the present limitations of the formulation, the hardware, and their current relationship to clinical practice and future translation.

Methods

Morphometric input

Parameters were obtained from our previously published CT-based study of 82 adult pelves (42 Turkish and 40 Russian) [5]. The primary instance used the Turkish male subgroup (n = 28).

Table 1
Morphometric inputs and candidate values for the primary Turkish male instance.

Parameter	Morphometric input	Source	Candidate values
Minimum corridor diameter	12.33 ± 2.04 mm	[5], Table 3	4.5, 6.5 mm
Maximum corridor length	125.39 ± 7.67 mm	[5], Table 3	115, 120, 125 mm
Craniocaudal angle	54.0° (IQR 4.50)	[5], Table 3	50°, 54°, 58°
Mediolateral angle	33.0° (IQR 6.00)	[5], Table 3	29°, 33°, 37°
Entry-point lateral offset	26.7 ± 3.8 mm	Ochs et al. [8]	24.0, 26.7, 30.0 mm

The entry-point offset was taken from Ochs et al. because the source study fixed the entry point at the medial aspect of the superior pubic ramus rather than treating it as a measured variable [5, 8]. Only published aggregate data and a publicly available de-identified pelvic model were used. No new patient contact, imaging, or measurements were performed. The public pelvic model was used only for the illustrative rendering shown in Fig. 1 [26].

Decision variables and candidate space

The Cartesian product yielded 162 configurations. The hardware benchmark used 115, 120, and 125 mm lengths, whereas cross-profile analysis used L − 10, L − 5, and L to preserve equal offsets from each profile-specific mean corridor length. These are model inputs, not claims about commercial screw inventory. For Turkish males, both constructions selected the same configuration type and differed only by rounding (cost 0.479 vs 0.473). Screw type and thread pattern were not modeled.

Constraints and cost function

The author, an orthopaedic surgeon, specified three hard exclusions:

1. 1. Screw diameter greater than the minimum corridor diameter.
2. 2. Screw length greater than the maximum corridor length.
3. 3. Cortical safety margin, defined as (corridor diameter - screw diameter)/2, below 1.0 mm.

Other clinically important conditions, including reduction quality, fracture-line orientation, early soft-tissue exit, and screw-head position, cannot be represented by this aggregate geometric parametrization and were not modeled.

The soft cost function was:

cost = 3.0|entry - entry_mean| + 2.0(1 - length/corridor_length) + 1.0(1 - diameter/corridor_diameter) + 0.5|sagittal - sagittal_mean| + 0.5|horizontal - horizontal_mean|

The weights combine clinical priority with unit scaling. Because every term is minimized simultaneously in this discretization, any strictly positive weighting selects the same configuration. Scale imbalance would matter only when anatomically coupled variables create genuine trade-offs; weight sensitivity was therefore tested directly.

QUBO encoding

The five lowest-cost configurations were encoded one-hot on five qubits, one qubit per candidate. Raw candidate costs were min-max normalized to [0,1], which preserves their ranking. The exactly-one

cardinality constraint was imposed as:

$QUBO = \Sigma_i c_i x_i + A(\Sigma_i x_i - 1)^2$

The penalty coefficient was A = 2.0. With normalized valid costs in [0,1], A > 1 places the all-zero state and other cardinality-violating states strictly above the complete valid-cost range. The QUBO was exhaustively enumerated over all 2^5 bitstrings, mapped to an Ising Hamiltonian using $x_i = (1 - z_i)/2$, and checked by exact diagonalization before hardware execution.

QAOA implementation and hardware execution

QAOA circuits were constructed at depths p = 1, 2, 3, and 4. Variational parameters were pre-optimized on a noiseless statevector simulator by minimizing expected QUBO energy with COBYLA, using 40 random restarts and up to 1500 iterations per restart. Success probability was not used as the optimization target because doing so would require prior knowledge of the solution.

Optimized circuits were transpiled at level 1 and executed with 1000 shots. The depth sweep used ibm_fez and ibm_marrakesh with four repeats per depth; the eight-profile and mitigation experiments used ibm_marrakesh. Hardware access was through a personal open/trial account, with approximately 3 s charged QPU time per job excluding queueing.

A Hadamard check produced 506/494 counts in simulation and 478/522 on ibm_marrakesh; neither differed from 50/50 at 1000 shots. This made gross circuit or measurement failure unlikely but did not quantify device noise. Each session also included an X-gate bit-order control; the expected bitstring was returned in 978/1000 shots in the eight-profile session.

Unconditional success probability was the fraction of all shots returning the optimal bitstring; its uniform baseline over 2^5 states is 1/32 = 0.031. Conditional success probability used only valid one-hot outcomes as the denominator; its baseline is 1/5 = 0.200.

On the eight-profile date, ibm_marrakesh median errors were 0.00286 for CZ, 0.00033 for single-qubit gates, and 0.01056 for readout; the job consumed 11 s charged QPU time. The depth-sweep estimate used same-day ibm_fez values (median CZ error 0.00282; readout error 0.00842).

Classical reference and software

Exhaustive enumeration of all 162 configurations provided ground truth. Simulated annealing (500 iterations with geometric cooling) and tabu-style steepest descent with random restarts were each evaluated from 100 independent starts. Analyses used Qiskit 2.5.1, Qiskit IBM Runtime 0.48.0, SciPy 1.16.3, and NumPy 2.0.2; COBYLA was used for variational optimization [22, 23]. The overall computational workflow is summarized in Fig. 2.

Declarations

Ethics statement

The study used only previously published aggregate morphometric data and a publicly available, de-identified pelvic model. No new patient-level data were collected, no patient contact or intervention occurred, and the public model was used only for illustrative rendering. Institutional review board approval and informed consent were therefore not required for this secondary analysis.

Use of artificial intelligence

Generative AI assistants (Anthropic Claude and OpenAI ChatGPT) were used to support code development, formulation checking, language editing, and manuscript drafting. The author designed the study, specified the clinical constraints and cost weights, executed and verified all computations, and takes full responsibility for the content and integrity of the work. AI tools were not used to generate experimental data and are not credited as authors.

Competing interests

The author declares no competing interests.

Funding

This study received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. Quantum hardware access was obtained through a personal open/trial account on the IBM Quantum Platform.

Author Contribution

Author ContributionsG.N.D. conceived and designed the study, defined the clinical constraints and optimization framework, performed the computational analyses and quantum hardware experiments, interpreted the results, prepared the figures and tables, and wrote the main manuscript text. G.N.D. reviewed and approved the final manuscript.

Data availability

All morphometric inputs and summary data supporting the findings are included in this article and the cited source studies. The public pelvic model used solely for Fig. 1 was derived from the CTPelvic1K resource [26]. De-identified hardware count data and optimized QAOA parameter sets are available from

the corresponding author on reasonable request and will be provided to editors and reviewers for assessment.

Code availability

The custom Python/Qiskit analysis code and cleaned notebooks used to generate the reported results are available from the corresponding author on reasonable request and will be provided to editors and reviewers for evaluation. No proprietary custom software was used.

References

1. Judet, R., Judet, J. & Letournel, E. Fractures of the acetabulum: classification and surgical approaches for open reduction. *J. Bone Joint Surg. Am.* **46**, 1615–1646 (1964).
2. Bozzio, A. E., Wydra, F. B., Mitchell, J. J., Ackerson, R. M. & Mauffrey, C. Percutaneous fixation of anterior and posterior column acetabular fractures. *Orthopedics* **37**, 675–678 (2014).
3. Gras, F. et al. Screw-versus plate-fixation strength of acetabular anterior column fractures. *J. Trauma. Acute Care Surg.* **72**, 1664–1670 (2012).
4. Mosheiff, R. & Liebergall, M. Maneuvering the retrograde medullary screw in pubic ramus fractures. *J. Orthop. Trauma.* **16**, 594–596 (2002).
5. Dirlik, G. N. et al. Morphometric analysis of the acetabular anterior column for retrograde screw fixation: a comparative CT study. *Indian J. Orthop.* https://doi.org/10.1007/s43465-026-01873-4 (2026).
6. Jarragh, A., Lari, A. & Shaikh, M. A CT-based morphometric study of superior pubic ramus anatomy among Arabs. *Surg. Radiol. Anat.* **45**, 603–609 (2023).
7. Chen, K. N., Wang, G., Cao, L. G. & Zhang, M. C. Differences in percutaneous retrograde screw fixation of anterior column acetabular fractures between males and females. *Injury* **40**, 1067–1072 (2009).
8. Ochs, B. G., Stuby, F. M., Ateschrang, A., Stoeckle, U. & Gonser, C. E. Retrograde lag screw placement in anterior acetabular column. *Injury* **45**, 1590–1598 (2014).
9. Trikha, V. et al. Morphometric analysis of the anterior column of the acetabulum and safety of intramedullary screw fixation in the Indian population. *Int. Orthop.* **44**, 655–664 (2020).
10. Dupuis, L. et al. Entry point variation in the osseous fixation pathway of the anterior column of the pelvis: a three-dimensional analysis. *J. Pers. Med.* **12**, 1748 (2022).
11. Attias, N. et al. Use of a virtual three-dimensional model to evaluate the intraosseous space available for percutaneous screw fixation of acetabular fractures. *J. Bone Joint Surg. Br.* **87**, 1520–1523 (2005).
12. Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. arXiv:1411.4028 (2014).

13. Lucas, A. Ising formulations of many NP problems. *Front. Phys.* **2**, 5. https://doi.org/10.3389/fphy.2014.00005 (2014).
14. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79. https://doi.org/10.22331/q-2018-08-06-79 (2018).
15. Jeyaraman, N., Jeyaraman, M. & Ramasubramanian, S. Quantum leap: envisioning the role of quantum computing in orthopedic practice. *J. Orthop. Case Rep.* **14**, 1–5 (2024).
16. Mallow, G. M. et al. Quantum computing: the future of big data and artificial intelligence in spine. *Spine Surg. Relat. Res.* **6**, 93–98. https://doi.org/10.22603/ssrr.2021-0251 (2022).
17. Mohamed, A. A. et al. Quantum computing in the realm of neurosurgery. *World Neurosurg.* **193**, 8–14. https://doi.org/10.1016/j.wneu.2024.09.131 (2025).
18. Khanna, R. et al. Quantum computing in surgery and urology – taking a quantum leap. *Nat. Rev. Urol.* **23**, 2–3. https://doi.org/10.1038/s41585-025-01058-y (2026).
19. Zhao, J. X. et al. Dependable automated approach for measuring the retrograde superior ramus screw corridor in pelvic fracture fixation. *J. Bone Joint Surg. Am.* **106**, 2356–2365 (2024).
20. Prasarn, M. L. et al. The effect of entry point on malalignment and iatrogenic fracture with the Synthes lateral entry femoral nail. *J. Orthop. Trauma.* **24**, 224–229. https://doi.org/10.1097/BOT.0b013e3181bedcc1 (2010).
21. Patel, A. H. et al. A review of proximal tibia entry points for intramedullary nailing and validation of the lateral parapatellar approach as extra-articular. *Orthop. Rev. (Pavia)*. **14**, 31909. https://doi.org/10.52965/001c.31909 (2022).
22. Javadi-Abhari, A. et al. Quantum computing with Qiskit. arXiv:2405.08810 (2024). https://doi.org/10.48550/arXiv.2405.08810
23. Powell, M. J. D. A direct search optimization method that models the objective and constraint functions by linear interpolation. In Advances in Optimization and Numerical Analysis (eds (eds Gomez, S. & Hennart, J. P.) 51–67 (Springer, Dordrecht, https://doi.org/10.1007/978-94-015-8330-5_4. (1994).
24. Zorrilla-Muñoz, V. et al. Springer, Cham,. Towards a framework for acetabular fracture and coxarthrosis risk prediction integrating quantum computing, machine learning, and 3D biomechanical modeling. In The 2026 International Conference on Breakthroughs in Artificial Intelligence (BAI'26) (eds Daimi, K. & Alsadoon, A.) Lecture Notes in Networks and Systems 1907 (2026). https://doi.org/10.1007/978-3-032-21147-7_11
25. Niroula, P. et al. Constrained quantum optimization for extractive summarization on a trapped-ion quantum computer. *Sci. Rep.* **12**, 17171. https://doi.org/10.1038/s41598-022-20853-w (2022).
26. Liu, P. et al. Deep learning to segment pelvic bones: large-scale CT datasets and baseline models. *Int. J. Comput. Assist. Radiol. Surg.* **16**, 749–756. https://doi.org/10.1007/s11548-021-02363-8 (2021).

Figures

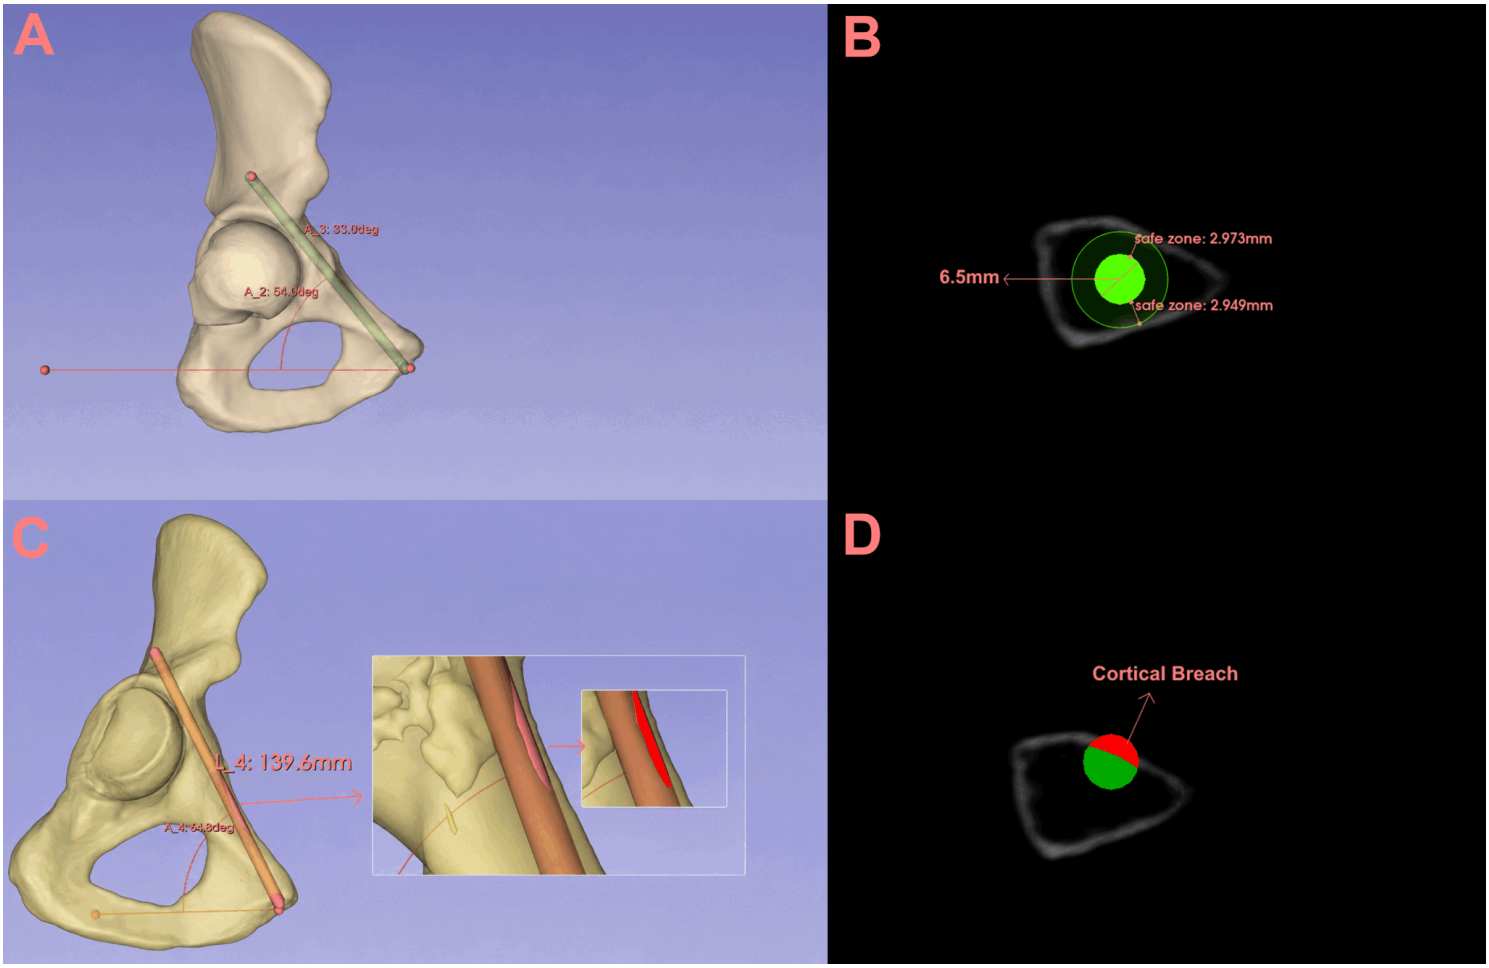


Figure 1

Illustrative trajectory placement on a publicly available CT-derived pelvic model. (a) The model-selected trajectory enters the superior pubic ramus at a craniocaudal angle of 54° and remains intraosseous. (b) Cross-sectional reformat perpendicular to the screw axis, showing an approximately 2.9-mm cortical margin around the 6.5-mm screw. (c) An exaggerated illustrative trajectory at approximately 64.8° breaches the cortex; this breach example was not one of the discretized candidate configurations. (d) Corresponding cross-sectional reformat. The figure is illustrative only and is not a clinical case, validation dataset, or demonstration of intended use. The model was derived from the publicly available CTPelvic1K dataset [26] and associated Zenodo release (https://doi.org/10.5281/zenodo.4588403) and modified for this study.

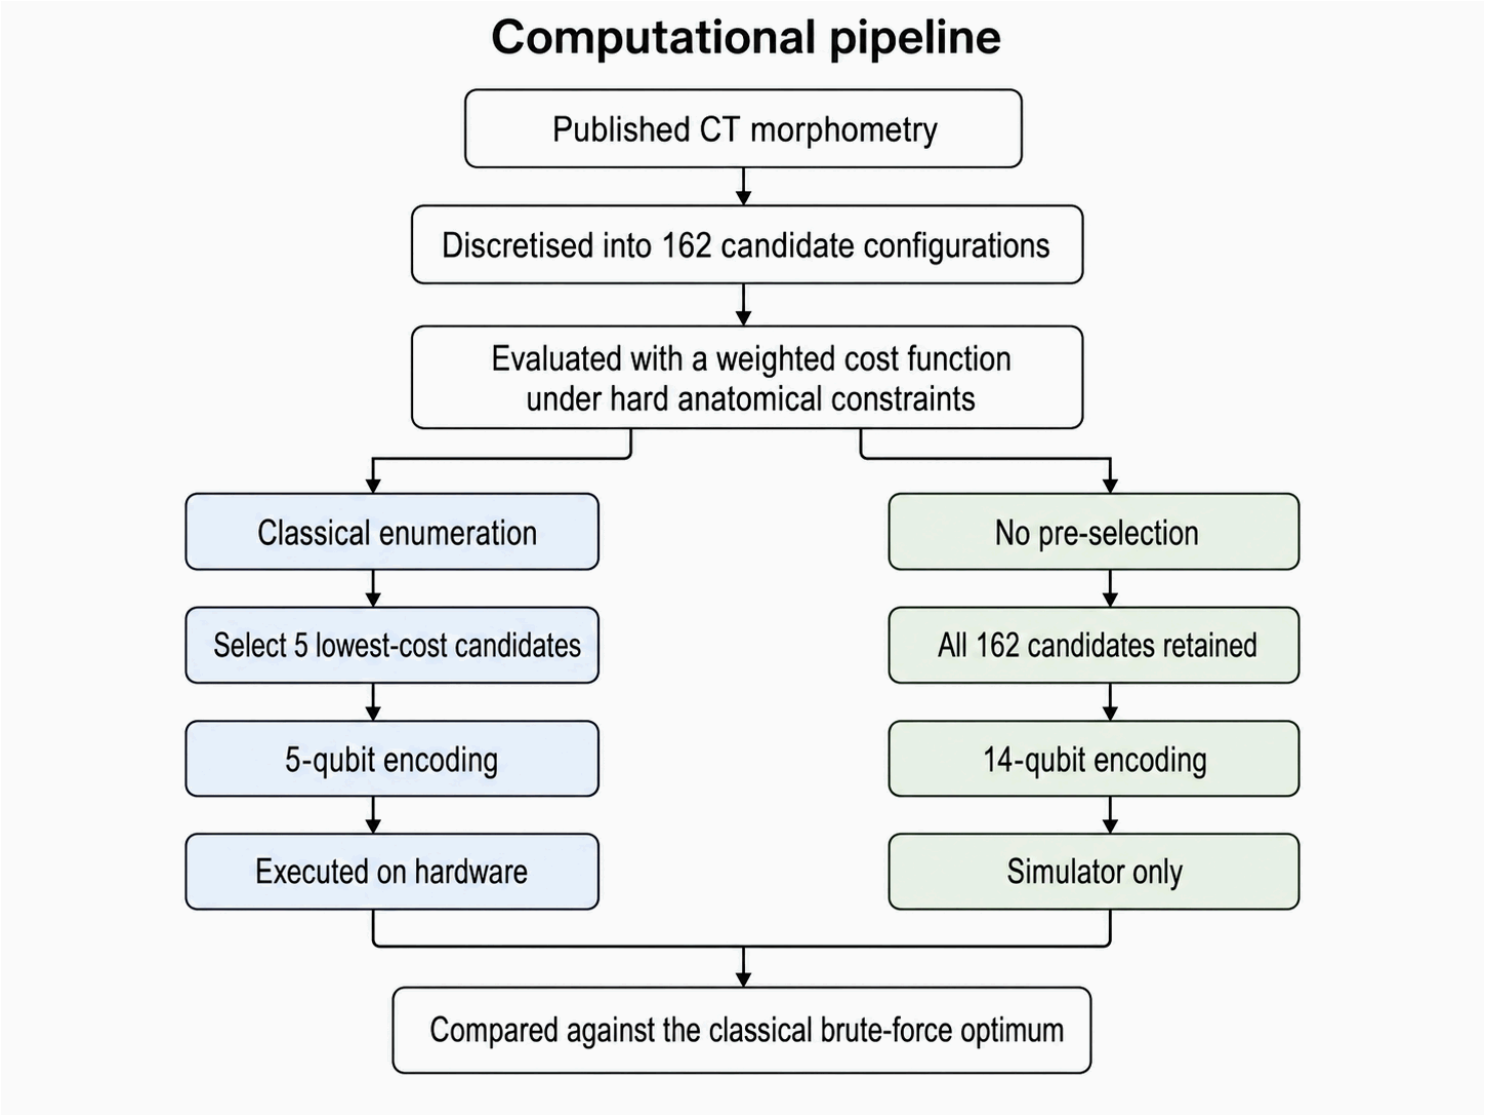


Figure 2

Computational pipeline. Published CT morphometry was discretized into 162 candidate configurations and evaluated using a weighted cost function and anatomical exclusions. The problem was encoded in two ways: a five-candidate shortlist selected by classical enumeration (5 qubits, executed on hardware) and the complete 162-configuration variable-wise formulation (14 qubits, simulator only). Both were compared with the classical exhaustive optimum.