C-QOBRA: QUANTUM CONDITIONING IN SPECIFIED MOLECULAR GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

C-QOBRA is a conditional variational autoencoder (cVAE) for molecular design, whose encoder and decoder are parameterized quantum circuits of eight to ten qubits with at most 204 trainable parameters. Read directly from the decoder circuit, with no classical network, it is already an effective *unconditional* generator: on QM9, it yields more unique-and-valid and more novel-and-valid molecules than fully quantum existing VAEs while using fewer quantum parameters, and most of its novel molecules stay within QM9's size limit. Conditional control is weaker under the same direct readout: requested logP values come out in the right order but are compressed to about a third of their range, and conditioning on QED barely moves the output. Because the circuit generates well, this points to decoding rather than conditioning as the bottleneck. Adding a classical post-processing head restores controllability and reveals how little classical capacity is needed: a head roughly $1/200$ the size of the largest tested already recovers most of the attainable logP control, with nearly five times fewer classical parameters than the pre-existing equivalent, which generates only unconditionally. Overall, the quantum circuits store the latent representation, and a small classical head suffices to decode it into conditioned property values.

1 INTRODUCTION

Generating molecules with specified physicochemical properties is a central problem in computer-aided drug discovery. A common approach trains a variational autoencoder (VAE) on molecules written as text (Gómez-Bombarelli et al., 2018): the encoder maps each string to a point in a continuous latent space, and new molecules are generated by decoding latent points steered toward target properties, so the text format bounds what the decoder can emit. SMILES (Weininger, 1987) writes a molecule as a string of atoms, bonds, and digits marking ring closures (e.g., `CCO` for ethanol); however, many strings describe no molecule. Alternatively, SELFIES (Krenn et al., 2020) is a grammar in which every string decodes to a molecular graph, i.e., atoms as nodes and bonds as edges.

Quantum circuits offer a compact alternative substrate for such generators (Benedetti et al., 2019; Cerezo et al., 2021). A circuit acts on a register of n_q *qubits*, whose joint state is a unit-norm vector of 2^{n_q} entries called *amplitudes*; eight qubits thus hold a 256-dimensional vector. The circuit is a fixed sequence of *gates*, each a norm-preserving linear map on that vector, and some gates rotate it by an adjustable angle. In a *parameterized quantum circuit* (PQC), these angles are the model parameters and are trained by gradient descent like network weights. Because the dimension of the vector grows exponentially in n_q while the parameter count grows only linearly, a PQC acts on a large space with few parameters. Its output, however, is accessible only through *measurement*, which returns averages over the state rather than the vector itself. It is therefore unclear how much of a molecular generation pipeline can be realized within a PQC and when classical components are required. Prior quantum generators (Li & Ghosh, 2022; Li et al., 2021; Gircha et al., 2023) use a classical decoder while rarely specifying its size; when they do, it outweighs the circuit by one to four orders of magnitude ($4{,}286$ classical vs. 108 quantum parameters in H-BQ-VAE, $453{,}644$ vs. 38 in QGAN-HG), so end-to-end metrics such as validity, uniqueness, and novelty cannot show how much of the performance the circuit itself supplies.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1.1 Related work

Classical conditional molecular generation. Beyond the string VAE (Kingma & Welling, 2022) described above, control is provided by the conditional VAE (Gulrajani et al., 2016) and the β-VAE objective (Higgins et al., 2017); the strongest generators on the QM9 benchmark are now equivariant diffusion models over 3-D coordinates (Hoogeboom et al., 2022; Xu et al., 2023), carrying millions of parameters.

Quantum generative models for molecules. Quantum autoencoders were introduced by Romero et al. (2017). Li & Ghosh (2022) proposes a baseline quantum VAE, whose hybrid form, H-BQ-VAE, adds a $4{,}286$-parameter classical output layer, and a scalable variant, SQ-VAE, that loads a flattened adjacency matrix into the amplitudes and decodes it classically; QGAN-HG (Li et al., 2021) pairs a hybrid quantum–classical generator with a classical discriminator that holds almost all of its parameters; and Gircha et al. (2023) replaces the latent restricted Boltzmann machine of a discrete VAE with a quantum annealer, a special-purpose quantum device that samples from such models. Closest to this work is QVAE-Mole (Wu et al., 2024), which generates 3-D point clouds from 224 quantum and no classical parameters. One of its impracticalities is that it reads latent code by tomography, which estimates every amplitude from many repeated measurements, a cost that grows exponentially with the number of qubits (Haah et al., 2016). Additionally, its conditional variant adds one extra qubit per property and reports the fraction of samples whose rounded property matches the request at a few isolated targets. This coarse form of evaluation does not allow for drawing deeper conclusions about conditional quantum VAEs. Generally, quantum and hybrid VAE outside of chemistry are unconditional and therefore complementary: ζ-QVAE (Wang et al., 2025) makes both encoder and decoder quantum, with latent codes stored as noisy quantum states known as mixed states, and VAE-QWGAN (Thomas et al., 2025) pairs a classical encoder with a quantum GAN generator.

What is and is not established. This line of work does show that a quantum circuit can generate valid, novel molecules, and that a conditional input can shift the output distribution. It does not show how much of either survives a parameter-free readout or how large a classical decoder must be before conditioning is recovered, because no study varies the capacity of the classical decoding stage. Those are the measurements we supply, in addition to building an architecture that is hardware economic and scale-efficient.

1.2 Contributions

We therefore characterize behavior along the quantum-classical axis. We present C-QOBRA, which extends QO-BRA (Yu et al., 2026) with the ability to embed conditions and generate molecules to a requested specification: a conditional VAE whose encoder and decoder are PQCs of eight or ten qubits with 160 or 204 trainable parameters, conditioned without adding qubits (Equation 1). The decoder output is read either (i) directly from the decoder circuit's measured output (*register readout*), with no trainable classical module between the circuit and the token sequence, or (ii) through a small Transformer post-processor. Register readout alone generates at least as well as QVAE-Mole (Section 4.2). However, like QVAE-Mole, preserves only the direction of a condition, not its scale (Section 4.3). Sweeping post-processor size over more than three orders of magnitude shows that most controllability returns within the first $\sim 10^3$ classical parameters (Section 4.4), still orders of magnitude below the classical generators mentioned above.

- We introduce conditioning without extra qubits: each property sets the angles of two-qubit rotations between adjacent existing qubits (Section 3.2), so the circuit maintains a fixed width regardless of how many properties it is conditioned on, whereas QVAE-Mole (Wu et al., 2024) adds one qubit per property, reaching eleven for its four-property experiment.
- Read directly from the decoder circuit with no classical network, the quantum core generates at least as well as the strongest fully quantum model: on QM9, $29.6 \pm 3.5\%$ of samples are unique and valid, and $62.3 \pm 4.6\%$ are novel and valid, compared to 27.4% and 57.4% for QVAE-Mole, with 204 rather than 224 quantum parameters (Table 1).
- We find that direct readout keeps the direction of a condition but not its scale: logP requests stay ordered, but their achieved range is compressed to about a third (slope 0.367; Figure 2a), a limitation that QVAE-Mole shares (Section 4.3). A classical head on a circuit

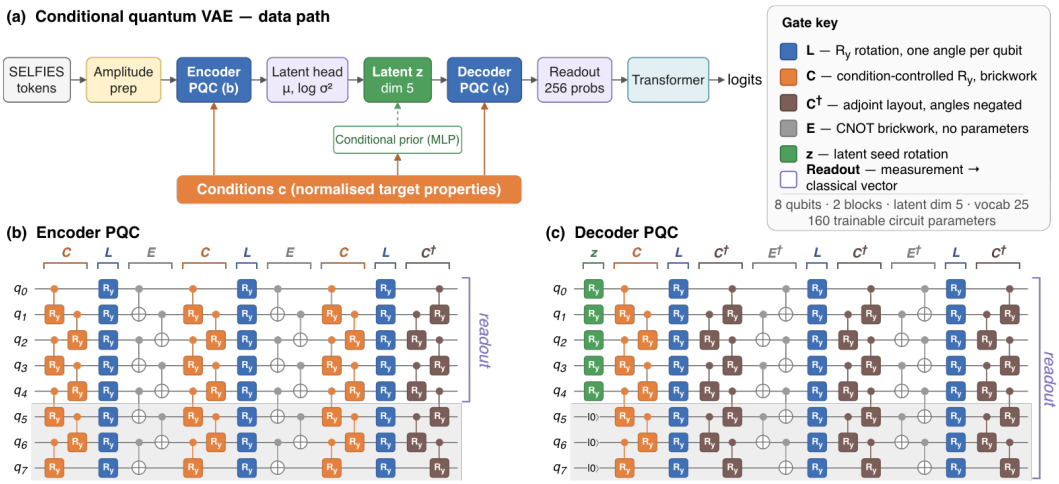

Figure 1: Overview of the C-QOBRA pipeline. **(a)** SELFIES token identifiers are zero-padded to 256 entries, ℓ_2-normalized, and amplitude-encoded into the encoder PQC **(b)**. The encoder outputs a 5-dimensional latent code that is passed to the decoder PQC **(c)**. The decoder produces $2^8 = 256$ output amplitudes, forming the memory representation for readout. Conditioning enters only through the controlled-rotation layers $C(\boldsymbol{c})$, four per circuit, with no extra qubits. The shown setup (eight qubits, two blocks, bottleneck dimension 3, one conditioning variable) has 160 trainable parameters. The register-readout arm uses the same construction at ten qubits (Section 3.5). Appendix A.2 details the full construction.

of the same design restores most of the scale, locating the bottleneck in decoding rather than conditioning.

- We measure the classical cost of that decoding by sweeping the head from 82 to 198,784 parameters (Figure 4): a 904-parameter head raises the logP slope to 0.742 ± 0.015, 84% of what the largest head reaches, and both it and a 3,344-parameter head are smaller than the 4,286 classical parameters that H-BQ-VAE (Li & Ghosh, 2022) uses for unconditional generation alone.

2 Preliminaries

Amplitude encoding writes a ℓ_2-normalized classical vector directly into the 2^{n_q} amplitudes, which is how classical data, such as a molecule, enters a circuit. The trainable gates of a PQC are typically single-qubit rotations $R_Y(\theta)$, and a particular arrangement of gates is called an *ansatz*, the circuit analog of a network architecture. Two-qubit gates such as CNOT create *entanglement*, i.e., correlations between qubits that no set of independent per-qubit states can reproduce; they play the role of the mixing layers in a network. Every gate is a real orthogonal matrix, so the adjoint $U^\dagger$ of a circuit U undoes it. A qubit can be measured along three axes, X, Y, and Z, each measurement returning $+1$ or -1; measuring qubit i along Z, and the standard axis gives the expectation value $\langle Z_i \rangle \in [-1, 1]$. Appendix A.1 states these definitions formally.

QO-BRA. C-QOBRA builds on QO-BRA (Yu et al., 2026), a quantum autoencoder for *de novo* design, where it generates from scratch instead of through library screening. It amplitude-encodes a token sequence, transforms it with an ansatz of R_Y rotations and CNOT gates between every pair of qubits, and decodes with the adjoint of the encoder, reusing its angles. Its latent code is the full 2^{n_q}-amplitude state, regularized toward a Gaussian, and it was evaluated on generating metal-binding peptides, i.e., short amino-acid chains, without embedding desired properties as conditions.

3 C-QOBRA Architecture

C-QOBRA is a cVAE whose encoder and decoder are PQCs. It generates SELFIES strings (Krenn et al., 2020) conditioned on continuous medicinal-chemistry descriptors, i.e., numerical properties

computed from a molecule's structure. All quantum layers are simulated in TorchQuantum (Wang et al., 2022). The circuits use no conditioning or extra qubits: every Transformer post-processor arm uses $n_q = 8$ ($d = 2^{n_q} = 256$ amplitudes), so comparisons across the post-processor sweep vary only the classical decoding head, while the register readout needs $n_q = 10$ (Section 3.5), so comparisons between the two readouts also differ in circuit width. Figure 1 traces the data flow; Appendix A.2 gives full specifications.

3.1 Encoding a molecule

Molecules are written as SELFIES strings and tokenized against a fixed 25-symbol vocabulary, with the alphabet occurring in the QM9 split rather than a SELFIES standard. Because every SELFIES string decodes to a molecular graph, this choice raises validity by design, favoring C-QOBRA in validity comparisons with models that output adjacency matrices or 3-D coordinates (Section 4.2); a sample still counts as invalid when RDKit rejects the decoded molecule. Token identifiers are embedded by amplitude encoding, where the vector is zero-padded to dimension d and then ℓ_2-normalized, with its entries becoming the amplitudes of the n_q-qubit state, so all $d = 256$ values are held by 8 qubits, and the preparation has no trainable parameters. Using raw token identifiers as amplitudes is a deliberate simplification that keeps the classical pre-encoding cost at exactly zero parameters.

3.2 Conditioning without extra qubits

Conditioning is where C-QOBRA departs most from earlier quantum molecular VAEs. As mentioned, QVAE-Mole (Wu et al., 2024) gives each property its own extra qubit, coupled to the data qubits by two-qubit rotations, so its circuit widens with every condition. We instead let the condition set gate angles on the existing qubits, which we arrange in a line: each of the m conditioning descriptors $c_k \in [-1, 1]$ enters through a layer of *controlled rotations* on the $n_q - 1$ pairs of adjacent qubits,

$$C(\boldsymbol{c}) \;=\; \prod_{k=1}^{m} \prod_{j=0}^{n_q-2} \mathrm{CRY}_{j,\,j+1}(w_{k,j}\, c_k + b_{k,j})\,, \tag{1}$$

where $\mathrm{CRY}_{j,j+1}$ applies the rotation to qubit $j{+}1$ only on the part of the state in which qubit j is 1, thereby entangling the pair, with all weights $w_{k,j}$ and biases $b_{k,j}$ being trainable by gradient descent exactly as in a classical layer, and with an affine map at each circuit's input and output. The condition, therefore, controls how strongly neighboring qubits become correlated, rather than merely rotating each qubit on its own, and re-applying $C(\boldsymbol{c})$ at every block, known as data re-uploading (Pérez-Salinas et al., 2020), feeds the condition in again at every depth, much as a conditional network concatenates the condition at every layer. The circuit remains at n_q qubits regardless of how many properties it is conditioned on. The decoder mirrors the encoder with a reversed layer order, $C^\dagger$ in place of C, and independently learned rotation angles.

3.3 Circuit structure

Each circuit alternates three kinds of layers: a trainable rotation layer (one R_Y per qubit), a parameter-free mixing layer of $n_q - 1 = 7$ CNOT gates between adjacent qubits, and the conditioning layer $C(\boldsymbol{c})$ of Equation 1. The encoder repeats rotation, mixing, and conditioning for n_{blocks} blocks and brackets the stack with conditioning layers, while the decoder mirrors it in reverse order with its own angles. Only the conditioning layers are specific to C-QOBRA: the backbone is the standard circuit-centric ansatz of Schuld et al. (2020), and because its two-qubit gates act only on adjacent qubits, it maps directly onto current hardware (Kandala et al., 2017). This compact mixing layer also distinguishes C-QOBRA from QO-BRA (Section 2), whose all-pairs layer needs $n_q(n_q - 1)/2 = 28$ CNOT gates for eight qubits: ours uses a quarter as many, yet still links each qubit to its neighbors so that correlations spread across the whole register over successive blocks, and it thereby limits two-qubit gates, a dominant source of error on current hardware (Preskill, 2018). Appendix A.2.3 writes both circuits out as equations.

3.4 LATENT SPACE AND PRIOR

Every gate in the stack (R_Y, CNOT, and CRY) has real matrix entries, and the amplitude-encoded input is a real vector, so the state remains real throughout. A qubit can be measured along the aforementioned three axes (X, Y, Z); for a real state, the Y average is always zero, so each of the five latent qubits carries a single real degree of freedom, which we read after the encoder as $\langle Z_i \rangle$. Because $\langle Z_i \rangle$ is bounded in $(-1, 1)$, we map it to the real line with $z_i = \operatorname{arctanh} \langle Z_i \rangle$, so the five latent qubits define an ordinary Gaussian latent in $\mathbb{R}^5$: the encoder's posterior is a diagonal Gaussian centered at these z_i, with variances from a small $\tanh$-activated head (Appendix A.2.4). This readout needs only one expectation value per latent qubit. For generation, the encoder is skipped, and latent codes are drawn from a learned conditional prior, a Gaussian whose mean and variance depend on $\boldsymbol{c}$, rather than a fixed standard normal, which places $\boldsymbol{z}$ in a specification-appropriate region from the start (Appendix A.2.5). The loss adds two quantum terms to the standard β-VAE objective of token-level reconstruction and a β-weighted KL regularizer with a free-bits threshold (Higgins et al., 2017; Kingma et al., 2016): the trash penalty of quantum autoencoders (Romero et al., 2017), which pushes the three non-latent trash qubits to a fixed state so that the input's information is compressed into the five latent qubits, and a purity penalty, introduced here, that keeps each latent qubit close to a pure (not mixed) state, where the $\operatorname{arctanh}$ readout stays informative (Appendix A.2.6). This qubit bottleneck, together with the one-number-per-qubit readout, distinguishes C-QOBRA's latent space from the full-state latent of QO-BRA (Section 2).

3.5 READING THE DECODER REGISTER

The state produced by the decoder circuit is a vector of 2^{n_q} amplitudes, which we call the *memory*. It holds the generated molecule in a distributed form rather than as tokens. How that memory is read is the axis that this paper varies, and we keep two readouts at its ends. The **register readout** partitions the register into a 5-qubit position sub-register and a 5-qubit token sub-register: measurement probabilities, marginalized over the remaining qubits, give $p(\text{position}, \text{token})$, and normalizing within each position gives $p(\text{token} \mid \text{position})$. Sequence position is thus a group of qubits rather than a network axis. Tokens are sampled independently per position, so this readout is non-autoregressive and adds no trainable parameters, classical or quantum; because the two sub-registers need 10 qubits, this arm uses $n_q = 10$ (Section 3.6). The **Transformer post-processor** instead reshapes the 256 amplitudes, without parameters, into $256/d_{\text{model}}$ memory vectors of width d_{model} and cross-attends to them with a one-layer Transformer decoder (Vaswani et al., 2017), using layer normalization before each sublayer (pre-LN; Xiong et al., 2020) rather than after it as in the original, while emitting tokens autoregressively, with an output head over the 25-token vocabulary at each step. Other examples of such hybrid quantum models would be H-BQ-VAE and SQ-VAE (Li & Ghosh, 2022), which pass the output of the circuit through a final classical linear layer. Neither the target tokens nor the memory vectors carry positional embeddings: token order comes from the causal mask on the target sequence alone (Haviv et al., 2022), and cross-attention treats the memory vectors as an unordered set. In Section 4.4, the only hyperparameter we vary is the width d_{model}, from 2 upward, with the feed-forward width tied to $2d_{\text{model}}$.

3.6 PARAMETER BUDGET

The quantum circuit has $2(n_{\text{blocks}}+1)n_q$ parametric rotation angles and $(n_{\text{blocks}}+1+\text{cond_depth}) \cdot 2m \cdot 2(n_q - 1)$ injection weights and biases. With $n_q = 8$, $n_{\text{blocks}} = 2$, $\text{cond_depth} = 1$, and a single conditioning descriptor ($m = 1$), this yields **160 trainable parameters**: 48 rotation and 112 injection. The 10-qubit register-readout circuit has, by the same formula, 204 (60 rotation and 144 injection). Other elements are non-parametric (e.g., amplitude encoding, register readout) or belong to the classical post-processor, whose size is varied as in Section 4.4; we therefore sweep post-processor sizes rather than fix a single end-to-end model.

Table 1: Unconditional generation on QM9, rates in %; all rows use 10,000 samples. Baseline values for QVAE-Mole, SQ-VAE (Li & Ghosh, 2022), and QGAN-HG and its P2 variant (Li et al., 2021) are as reported by Wu et al. (2024); (C-)QOBRA denotes C-QOBRA sampled without its conditional inputs, as mean ± sd over five seeds. QO-BRA (Yu et al., 2026) is not listed because it was evaluated on peptides rather than QM9. Following Wu et al. (2024), Unique×Valid and Novel×Valid are the fractions of all samples that are valid and unique, or valid and novel. #C/Q counts classical and quantum network parameters.

Method	Output	Valid (%)	Unique×Valid (%)	Novel×Valid (%)	#C/Q
QGAN-HG	graph	66.64 ± 0.3	8.08 ± 0.8	18.48 ± 1.0	453,644 / 38
P2-QGAN-HG	graph	17.64 ± 1.2	12.38 ± 2.1	9.54 ± 0.8	112,524 / 14
SQ-VAE	graph	44.23 ± 1.0	7.24 ± 1.2	16.32 ± 0.8	128 / 224
QVAE-Mole	point cloud	78.13 ± 0.6	27.39 ± 0.8	57.38 ± 0.9	0 / 224
(C-)QOBRA (ours)	SELFIES	$\mathbf{86.8 \pm 6.8}$	$\mathbf{29.6 \pm 3.5}$	$\mathbf{62.3 \pm 4.6}$	**0 / 204**

4 RESULTS

4.1 EXPERIMENTAL SETUP

All models were trained on QM9 (Glavatskikh et al., 2019), a benchmark of small organic molecules with at most nine heavy (non-hydrogen) atoms that is widely used in molecular machine learning (Glavatskikh et al., 2019; Pinheiro et al., 2020; Wu et al., 2024). Descriptors were computed with the RDKit cheminformatics toolkit (Landrum et al., 2026) and were affinely rescaled to $[-1, 1]$ so that user inputs map directly to a rotation angle (Section A.2.2), and the latent dimensionality was fixed at 5. Our primary conditioning target is logP (the log octanol–water partition coefficient, a measure of hydrophobicity), estimated with RDKit's Crippen implementation (Wildman & Crippen, 1999), with seven targets from -3 to $+3$ (step 1) and 1,000 molecules per target.

Fsp3 (the fraction of carbons that are sp^3-hybridized, i.e., bonded by four single bonds), a medicinal-chemistry–relevant descriptor (Wei et al., 2020), is the supporting case, with six targets from 0.0 to 1.0 in steps of 0.2 and 1,000 molecules per target. Further, we consider QED (quantitative estimate of drug-likeness, a single score in $[0, 1]$ aggregating eight molecular properties (Bickerton et al., 2012)) and fScafHetero (the fraction of heteroatoms, i.e., non-carbon heavy atoms, in the scaffold, the ring system left after removing side chains).

For Section 4.2, we drew 10,000 unconditional samples from the prior for each of the five training seeds, with no conditioning or post-processing; all samples were evaluated. Validity was determined using RDKit, and novelty was defined as the absence from the training set. These two, along with uniqueness, follow standard molecular-generation benchmarks (Polykovskiy et al., 2020; Brown et al., 2019). We report the product of validity with uniqueness and novelty, respectively, in the format of Wu et al. (2024).

Steering is summarized by: (i) the regression slope of the median achieved value versus the requested targets, (ii) R^2, and (iii) the mean absolute percentage error (MAPE) of all generated values versus the target, excluding target 0. We also report the span of achieved medians to reflect the effect size that correlation can miss (e.g., medians 0.50–0.60 can still correlate well but yield limited steering). Each post-processor size, or **arm**, in the sweep is trained under five seeds and reported as the mean ± sample standard deviation.

4.2 GENERATION

Before evaluation steering, we establish that it generates, in the most stringent setting available: direct register readout of the ten-qubit circuit (Section 3.5), with no classical network between the decoder circuit and the emitted string. Table 1 compares this against four prior quantum and hybrid models, including QVAE-Mole (Wu et al., 2024), the only fully quantum one, using its own metrics and sample counts. Because SELFIES raises validity by design (Section 3.1), the comparison rests on Unique×Valid and Novel×Valid, which it does not guarantee: C-QOBRA is higher on both (29.6 vs. 27.4% and 62.3 vs. 57.4%) with 204 rather than 224 quantum parameters, though on ten qubits rather than seven. This is consistent with our architecture generating at least as well with fewer

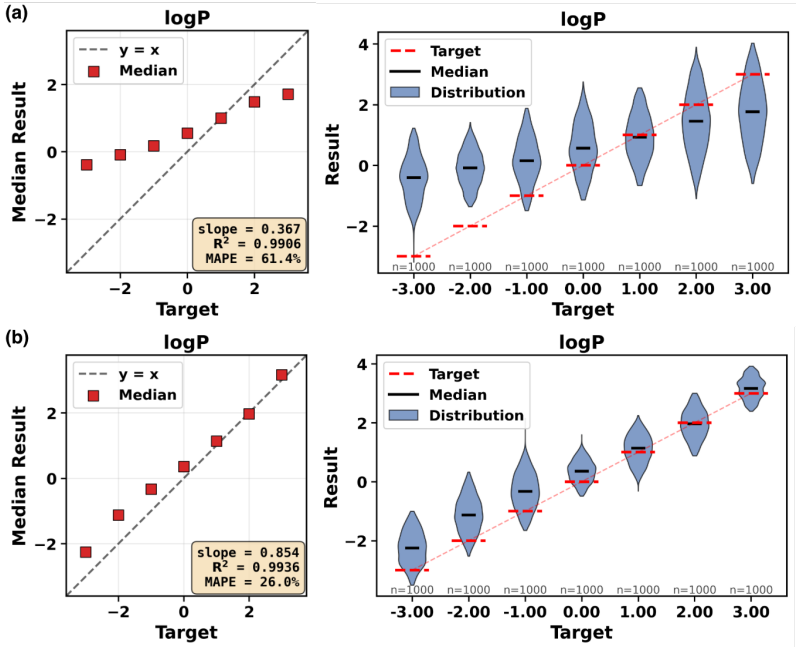

Figure 2: logP-conditioned generation (1,000 molecules for each of seven targets from -3 to $+3$); each left panel prints its fit. **(a)** Register readout without post-processing: medians rise monotonically and are well ordered but cover only about a third of the requested range. **(b)** The same circuit design, at eight rather than ten qubits, with a $d_{\text{model}} = 128$ post-processor, initialized with the same seed. Medians approach the diagonal and per-target distributions tighten, with a residual shortfall at the most extreme requests. The five-seed mean slope is 0.879 ± 0.028 (Table 2).

trainable parameters, although both margins lie within our seed spread. The hybrid graph generators in their comparison fall well below both models on every rate. The quantum core thus generates at least as well as the strongest fully quantum model before any classical decoder is added. Table 1 covers quantum and hybrid models only; a classical VAE is not parameter-matched to this regime, its counterpart in Li & Ghosh (2022) carrying 5,694 trainable parameters against the 204 here, and the strongest QM9 generators carrying millions (Section 1.1). That study compares the two families directly across its experiments.

Novelty could also be inflated by sampling molecules larger than any seen in training, but $85.6 \pm 5.6\%$ of the novel set respects QM9's nine-heavy-atom limit. Two such samples, the aromatic heterocycle `OCCc1ccncn1` (a ring with a non-carbon atom and alternating double bonds) and the fully saturated (single-bonded) `CCC1CNCCOC1`, sit at that bound at opposite ends of the saturation axis, suggesting the sampler does not collapse onto a narrow region of chemical space.

4.3 Readout

Conditioning is injected at eight sites across the two circuits and into the learned prior, but injection during training does not guarantee that the condition survives decoding to a molecule. Under register readout, the logP signal is present but miscalibrated (Figure 2a): medians rise monotonically with the request, and the fit is well ordered; yet, with slope 0.367, they span only about a third of the requested six-unit range, and the per-target distributions overlap heavily. Analogously, QVAE-Mole, which also decodes without classical parameters, displays its own shortcomings in conditional generation. In its Figure 5 multi-condition experiment, a logP request of 0.92 moves the median of the generated molecules from about -0.5 without conditioning to only about -0.2, roughly a quarter of the requested shift, while the distribution still spans several logP units. At the nearest request in our grid, logP $= 1$, the register readout alone places the median at about 1.0 (Figure 2a). In both models, a fully quantum readout carries the direction of the condition but not its full scale.

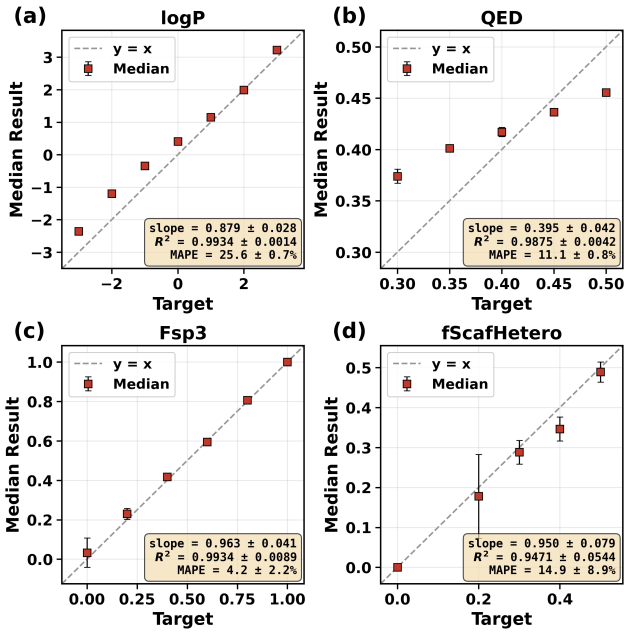

Figure 3: Median achieved value against the requested target for **(a)** logP, **(b)** QED, **(c)** Fsp3, and **(d)** fScafHetero, each decoded by the same 198,784-parameter ($d_{\text{model}} = 128$) post-processor; points are means and bars sample standard deviations over five seeds, and each panel prints its fit. Three of the four track the request closely, fScafHetero with the widest seed spread; QED remains weakly steered even at this capacity. Target grids differ by descriptor, so MAPE is not comparable across panels.

Adding a $d_{\text{model}} = 128$ post-processor to our circuits closes most of this gap: the slope rises to 0.879 ± 0.028 over five seeds, and the per-target distributions narrow (Figure 2b). For logP, the head therefore amplifies and calibrates a signal that is already there, rather than restoring an absent one. As illustrated in Figure 3, the same head steers the other descriptors as well. Fsp3 and fScafHetero track the request closely, whereas QED remains weakly steered even at this capacity (A.3, A.5, & A.6). This is likely because the aforementioned aggregation gives QED a relationship to molecular constituents that C-QOBRA does not capture well. The remaining question is what minimal, principled form this post-processing should take.

4.4 Post-Processor

Figure 4 sweeps the post-processor for all four descriptors across more than three orders of magnitude, and three shapes appear. For logP and Fsp3, slope rises and MAPE falls steeply over the first $\sim 10^3$ parameters and then flattens, so most of the attainable steering is obtained early; fScafHetero improves more gradually and non-monotonically; and QED never becomes steerable at any capacity.

Following logP through the sweep, a 904-parameter head ($d_{\text{model}} = 8$, 0.5% of the largest) already raises the logP slope from the register readout's 0.367 to 0.742 ± 0.015, 84% of what the largest head reaches. For comparison, H-BQ-VAE (Section 1.1) attaches $4{,}286$ classical parameters to its 108-parameter circuit and generates only unconditionally; both our 904- and $3{,}344$-parameter heads are smaller, yet they support conditional generation. Beyond $\sim 10^3$ parameters, the gain is slow: 220 times more parameters add only 0.14 to the slope (Table 2). Fsp3 shows the same shape more sharply (Figure 4c; Appendix A.3).

Steering saturates while reconstruction does not: token-level reconstruction accuracy rises monotonically with head size for every descriptor, including over the range where steering is flat (Table 2). Reconstruction requires the head to reproduce the tokens of an observed molecule, whereas steering only asks it to render a prior latent that already carries the condition, so added capacity keeps helping the former after the latter saturates.

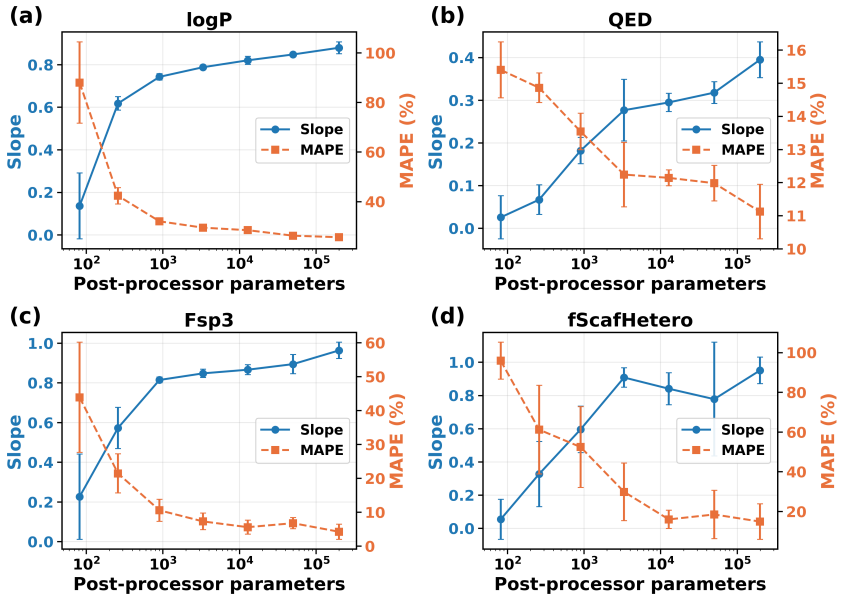

Figure 4: Steering as a function of post-processor size for **(a)** logP, **(b)** QED, **(c)** Fsp3, and **(d)** fScafHetero: slope (left axis) and MAPE (right axis) of the median achieved values (Section 4.1). Circuits, conditional prior and training objective are fixed; only the decode head varies, from $d_{\text{model}} = 2$ to 128 (82 to 198,784 parameters). Points are means and bars sample standard deviations over five seeds. MAPE is comparable within a panel but not across panels; per-arm values are in Table 2.

A classical module is not sufficient by itself, however: at 82 parameters, the logP slope falls to 0.136 ± 0.155, below even the register readout, with a seed spread larger than its mean, and on Fsp3, one seed collapses to constant medians (Table 2; Appendix A.4 explains why we report slope and MAPE rather than Pearson correlation).

5 Discussion and Conclusion

Read directly from the decoder circuit, with no classical network, C-QOBRA generates as effectively as the strongest fully quantum molecular VAE, QVAE-Mole, on unique-valid and novel-valid rates, with fewer quantum parameters (Table 1). The same direct readout carries the direction of a condition but not its scale (Section 4.3), a limitation that QVAE-Mole shares. Attaching a small classical head to a register from the same circuit design restores most of the scale, so the condition is present in that register, and the head supplies only the mapping from register state to token sequence; as that circuit is wider and trained separately, this compares architectures rather than readouts on fixed weights.

The largest head carries 198,784 trainable parameters against 160 in the circuit, above 1,200:1. The sweep supplies the missing intermediate regime: a 904-parameter head, smaller than the 4,286 classical parameters of H-BQ-VAE, already recovers most of the attainable steering, and gains beyond it are slow. Hence, representational capacity sits in the circuit, while conditional decoding occurs in the head.

Two limits qualify these results: the generation comparison crosses output representations, as benchmarks commonly do (Polykovskiy et al., 2020; Brown et al., 2019), with SELFIES favoring validity (Krenn et al., 2020; 2022), and QED is not steered at any head size we tested. The steering is achieved only with a classical module, so the pipeline is no longer fully quantum; whether the register could instead be read out by a quantum circuit rather than a classical head is the natural next question, alongside a parameter-matched classical control and a chemical space beyond QM9.

AI USE STATEMENT

We used generative AI tools in two roles, both under the authors' supervision.

Code: The implementation of the hybrid quantum–classical model was written with the assistance of Anthropic's Claude Code. The authors specified the architecture, the training objective, and the evaluation metrics; the assistant produced and iterated on the implementing code under the authors' direction. All code was read, executed, and checked by the authors, and every number reported here was produced by running that code.

Manuscript: Drafting and revision of this manuscript were assisted by Anthropic's Claude, used through the Claude in Chrome browser integration to edit the source directly. The authors determined the claims, the framing, and the scope of every result; the assistant contributed to drafting, restructuring, copy-editing, and flagged inconsistencies between reported numbers and the underlying run logs, which the authors then verified against the experiment outputs.

We did not use generative AI to generate, alter, or interpolate any experimental data, and no reported number was produced by a language model. We have reviewed all AI-assisted work and take full responsibility for the final content of this paper, including all text, claims, and artifacts.

REPRODUCIBILITY STATEMENT

The C-QOBRA code will be released upon acceptance and will fully reproduce all presented results.

ACKNOWLEDGMENTS

Acknowledgments will be made upon acceptance.

REFERENCES

Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001, November 2019. ISSN 2058-9565. doi: 10.1088/2058-9565/ab4eb5. URL https://doi.org/10.1088/2058-9565/ab4eb5.

G. Richard Bickerton, Gaia V. Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L. Hopkins. Quantifying the chemical beauty of drugs. *Nature Chemistry*, 4(2):90–98, February 2012. ISSN 1755-4349. doi: 10.1038/nchem.1243. URL https://www.nature.com/articles/nchem.1243.

Nathan Brown, Marco Fiscato, Marwin H.S. Segler, and Alain C. Vaucher. GuacaMol: Benchmarking Models for de Novo Molecular Design. *Journal of Chemical Information and Modeling*, 59(3):1096–1108, March 2019. ISSN 1549-9596. doi: 10.1021/acs.jcim.8b00839. URL https://doi.org/10.1021/acs.jcim.8b00839.

M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. Variational quantum algorithms. *Nature Reviews Physics*, 3(9):625–644, September 2021. ISSN 2522-5820. doi: 10.1038/s42254-021-00348-9. URL https://www.nature.com/articles/s42254-021-00348-9.

A. I. Gircha, A. S. Boev, K. Avchaciov, P. O. Fedichev, and A. K. Fedorov. Hybrid quantum-classical machine learning for generative chemistry and drug design. *Scientific Reports*, 13(1):8250, May 2023. ISSN 2045-2322. doi: 10.1038/s41598-023-32703-4. URL https://www.nature.com/articles/s41598-023-32703-4.

Marta Glavatskikh, Jules Leguy, Gilles Hunault, Thomas Cauchy, and Benoit Da Mota. Dataset's chemical diversity limits the generalizability of machine learning predictions. *Journal of Cheminformatics*, 11(1):69, November 2019. ISSN 1758-2946. doi: 10.1186/s13321-019-0391-2. URL https://doi.org/10.1186/s13321-019-0391-2.

Ishaan Gulrajani, Kundan Kumar, Faruk Ahmed, Adrien Ali Taiga, Francesco Visin, David Vazquez, and Aaron Courville. PixelVAE: A Latent Variable Model for Natural Images, November 2016. URL http://arxiv.org/abs/1611.05013. arXiv:1611.05013 [cs.LG].

Rafael Gómez-Bombarelli, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science*, 4(2):268–276, February 2018. ISSN 2374-7943. doi: 10.1021/acscentsci.7b00572. URL https://doi.org/10.1021/acscentsci.7b00572.

Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. Sample-optimal tomography of quantum states. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, STOC '16, pp. 913–925, New York, NY, USA, June 2016. Association for Computing Machinery. ISBN 978-1-4503-4132-5. doi: 10.1145/2897518.2897585. URL https://dl.acm.org/doi/10.1145/2897518.2897585.

Adi Haviv, Ori Ram, Ofir Press, Peter Izsak, and Omer Levy. Transformer Language Models without Positional Encodings Still Learn Positional Information. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 1382–1390, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-emnlp.99. URL https://aclanthology.org/2022.findings-emnlp.99/.

Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. February 2017. URL https://openreview.net/forum?id=Sy2fzU9gl.

Emiel Hoogeboom, Vıctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant Diffusion for Molecule Generation in 3D. In *Proceedings of the 39th International Conference on Machine Learning*, pp. 8867–8887. PMLR, June 2022. URL https://proceedings.mlr.press/v162/hoogeboom22a.html.

Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, September 2017. ISSN 1476-4687. doi: 10.1038/nature23879. URL https://www.nature.com/articles/nature23879.

Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes, December 2022. URL http://arxiv.org/abs/1312.6114. arXiv:1312.6114 [stat.ML].

Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved Variational Inference with Inverse Autoregressive Flow. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper_files/paper/2016/hash/ddeebdeefdb7e7e7a697e1c3e3d8ef54-Abstract.html.

Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alan Aspuru-Guzik. Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4):045024, October 2020. ISSN 2632-2153. doi: 10.1088/2632-2153/aba947. URL https://doi.org/10.1088/2632-2153/aba947.

Mario Krenn, Qianxiang Ai, Senja Barthel, Nessa Carson, Angelo Frei, Nathan C. Frey, Pascal Friederich, Théophile Gaudin, Alberto Alexander Gayle, Kevin Maik Jablonka, Rafael F. Lameiro, Dominik Lemm, Alston Lo, Seyed Mohamad Moosavi, José Manuel Nápoles-Duarte, AkshatKumar Nigam, Robert Pollice, Kohulan Rajan, Ulrich Schatzschneider, Philippe Schwaller, Marta Skreta, Berend Smit, Felix Strieth-Kalthoff, Chong Sun, Gary Tom, Guido Falk von Rudorff, Andrew Wang, Andrew D. White, Adamo Young, Rose Yu, and Alán Aspuru-Guzik. SELFIES and the future of molecular string representations. *Patterns*, 3(10), October 2022. ISSN 2666-3899. doi: 10.1016/j.patter.2022.100588. URL https://www.cell.com/patterns/abstract/S2666-3899(22)00206-9.

Greg Landrum, Paolo Tosco, Brian Kelley, Ricardo Rodriguez, David Cosgrove, Riccardo Vianello, sriniker, Peter Gedeck, Gareth Jones, Eisuke Kawashima, NadineSchneider, Dan Nealschneider, tadhurst-cdd, Andrew Dalke, Matt Swain, Brian Cole, Samo Turk, Aleksandr Savelev, Niels Maeder, Rachel Walker, Alain Vaucher, Maciej Wójcikowski, Hussein Faara, Ichiru Take, Vincent F. Scalfani, Yakov Pechersky, Kazuya Ujihara, Daniel Probst, Jeremy Monat, and Juuso Lehtivarjo. rdkit/rdkit: 2026_03_1 (Q1 2026) Release. *Zenodo*, March 2026. doi: 10.5281/zenodo.19250388. URL https://ui.adsabs.harvard.edu/abs/2026zndo..19250388L. ADS Bibcode: 2026zndo..19250388L.

Junde Li and Swaroop Ghosh. Scalable Variational Quantum Circuits for Autoencoder-based Drug Discovery. In *2022 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 340–345, March 2022. doi: 10.23919/DATE54114.2022.9774564. URL https://ieeexplore.ieee.org/abstract/document/9774564. ISSN: 1558-1101.

Junde Li, Rasit O. Topaloglu, and Swaroop Ghosh. Quantum Generative Models for Small Molecule Drug Discovery. *IEEE Transactions on Quantum Engineering*, 2:1–8, 2021. ISSN 2689-1808. doi: 10.1109/TQE.2021.3104804. URL https://ieeexplore.ieee.org/abstract/document/9520764.

Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812, November 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-07090-4. URL https://www.nature.com/articles/s41467-018-07090-4.

Mikko Mottonen, Juha J. Vartiainen, Ville Bergholm, and Martti M. Salomaa. Transformation of quantum states using uniformly controlled rotations, July 2004. URL http://arxiv.org/abs/quant-ph/0407010. arXiv:quant-ph/0407010.

Gabriel A. Pinheiro, Johnatan Mucelini, Marinalva D. Soares, Ronaldo C. Prati, Juarez L. F. Da Silva, and Marcos G. Quiles. Machine Learning Prediction of Nine Molecular Properties Based on the SMILES Representation of the QM9 Quantum-Chemistry Dataset. *The Journal of Physical Chemistry A*, 124(47):9854–9866, November 2020. ISSN 1089-5639. doi: 10.1021/acs.jpca.0c05969. URL https://doi.org/10.1021/acs.jpca.0c05969.

Daniil Polykovskiy, Alexander Zhebrak, Benjamin Sanchez-Lengeling, Sergey Golovanov, Oktai Tatanov, Stanislav Belyaev, Rauf Kurbanov, Aleksey Artamonov, Vladimir Aladinskiy, Mark Veselov, Artur Kadurin, Simon Johansson, Hongming Chen, Sergey Nikolenko, Alán Aspuru-Guzik, and Alex Zhavoronkov. Molecular Sets (MOSES): A Benchmarking Platform for Molecular Generation Models. *Frontiers in Pharmacology*, 11, December 2020. ISSN 1663-9812. doi: 10.3389/fphar.2020.565644. URL https://www.frontiersin.org/journals/pharmacology/articles/10.3389/fphar.2020.565644/full.

John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018. doi: 10.22331/q-2018-08-06-79. URL https://quantum-journal.org/papers/q-2018-08-06-79/.

Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I. Latorre. Data re-uploading for a universal quantum classifier. *Quantum*, 4:226, February 2020. doi: 10.22331/q-2020-02-06-226. URL https://quantum-journal.org/papers/q-2020-02-06-226/.

Jonathan Romero, Jonathan P Olson, and Alan Aspuru-Guzik. Quantum autoencoders for efficient compression of quantum data. *Quantum Science and Technology*, 2(4):045001, August 2017. ISSN 2058-9565. doi: 10.1088/2058-9565/aa8072. URL https://doi.org/10.1088/2058-9565/aa8072.

Maria Schuld and Nathan Killoran. Quantum Machine Learning in Feature Hilbert Spaces. *Physical Review Letters*, 122(4):040504, February 2019. doi: 10.1103/PhysRevLett.122.040504. URL https://link.aps.org/doi/10.1103/PhysRevLett.122.040504.

Maria Schuld, Alex Bocharov, Krysta M. Svore, and Nathan Wiebe. Circuit-centric quantum classifiers. *Physical Review A*, 101(3):032308, March 2020. doi: 10.1103/PhysRevA.101.032308. URL https://link.aps.org/doi/10.1103/PhysRevA.101.032308.

Aaron Mark Thomas, Harry Youel, and Sharu Theresa Jose. VAE-QWGAN: addressing mode collapse in quantum GANs via autoencoding priors. *Quantum Machine Intelligence*, 7(2):91, September 2025. ISSN 2524-4914. doi: 10.1007/s42484-025-00314-z. URL https://doi.org/10.1007/s42484-025-00314-z.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

Gaoyuan Wang, Jonathan Warrell, Prashant S. Emani, and Mark Gerstein. Quantum variational autoencoder utilizing regularized mixed-state latent representations. *Physical Review A*, 111(4): 042416, April 2025. doi: 10.1103/PhysRevA.111.042416. URL https://link.aps.org/doi/10.1103/PhysRevA.111.042416.

Hanrui Wang, Yongshan Ding, Jiaqi Gu, Yujun Lin, David Z. Pan, Frederic T. Chong, and Song Han. QuantumNAS: Noise-Adaptive Search for Robust Quantum Circuits. In *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 692–708, April 2022. doi: 10.1109/HPCA53966.2022.00057. URL https://ieeexplore.ieee.org/abstract/document/9773233. ISSN: 2378-203X.

Wenxiu Wei, Srinivasulu Cherukupalli, Lanlan Jing, Xinyong Liu, and Peng Zhan. Fsp3: A new parameter for drug-likeness. *Drug Discovery Today*, 25(10):1839–1845, October 2020. ISSN 1359-6446. doi: 10.1016/j.drudis.2020.07.017. URL https://www.sciencedirect.com/science/article/pii/S135964462030297X.

David Weininger. SMILES, a Chemical Language and Information System. 1. Introduction to Methodology and Encoding Rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, June 1987. doi: 10.1021/ci00057a005.

Scott A. Wildman and Gordon M. Crippen. Prediction of Physicochemical Parameters by Atomic Contributions. *Journal of Chemical Information and Computer Sciences*, 39(5):868–873, September 1999. ISSN 0095-2338. doi: 10.1021/ci990307l. URL https://doi.org/10.1021/ci990307l.

Huanjin Wu, Xinyu Ye, and Junchi Yan. QVAE-Mole: The Quantum VAE with Spherical Latent Variable Learning for 3-D Molecule Generation. In *Advances in Neural Information Processing Systems*, volume 37, pp. 22745–22771. Curran Associates, Inc., 2024. doi: 10.52202/079017-0716. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/286e7ab0ce6a68282394c92361c27b57-Abstract-Conference.html.

Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. On Layer Normalization in the Transformer Architecture. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 10524–10533. PMLR, November 2020. URL https://proceedings.mlr.press/v119/xiong20b.html.

Minkai Xu, Alexander S. Powers, Ron O. Dror, Stefano Ermon, and Jure Leskovec. Geometric Latent Diffusion Models for 3D Molecule Generation. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 38592–38610. PMLR, July 2023. URL https://proceedings.mlr.press/v202/xu23n.html.

Yue Yu, Francesco Calcagno, Haote Li, and Victor S. Batista. QO-BRA: A Quantum Operator-Based Autoencoder for De Novo Molecular Design. *Journal of Chemical Theory and Computation*, 22 (4):2059–2073, February 2026. ISSN 1549-9618. doi: 10.1021/acs.jctc.5c01704. URL https://doi.org/10.1021/acs.jctc.5c01704.

A Appendix

A.1 Preliminaries

Let $\mathcal{M}$ denote a set of molecular graphs, $\mathcal{A}$ a finite token alphabet, and $\mathcal{A}^{\leq T}$ the set of token sequences of length at most T. A **molecular string representation** is a mapping $s : \mathcal{M} \to \mathcal{A}^{\leq T}$. The aforementioned SMILES (Weininger, 1987) is a widely used representation. However, the induced decoding map is not surjective onto the set of chemically valid molecules, since many token sequences violate valence or ring-closure constraints; hence, they do not decode to any molecule. By contrast, SELFIES (Krenn et al., 2020), introduced in Section 1, is designed such that every token sequence maps to a syntactically valid molecule. We refer to this guarantee as the **closure property**, which enables generative models defined over token sequences to ensure validity by construction, motivating the use of SELFIES rather than SMILES in C-QOBRA.

Let $\boldsymbol{c} : \mathcal{M} \to \mathbb{R}^m$ denote a vector-valued descriptor function, and define $\boldsymbol{c}(x)$ as the **specification** of a molecule x. Descriptors are computed using RDKit (Landrum et al., 2026) and then affinely rescaled to the interval $[-1, 1]$ prior to use, so that specifications may be interpreted directly as rotation angles (Section A.2.2). A **conditional generative model** is a parameterized family $p_\theta(x \mid \boldsymbol{c})$ that supports sampling molecules conditioned on a target specification; the primary evaluation criterion is the fidelity with which the realized descriptor values of generated samples match the requested specification.

The cVAE underlying C-QOBRA (Section 3) specifies $p_\theta(x \mid \boldsymbol{c})$ by introducing a latent variable $\boldsymbol{z}$ with a prior $p(\boldsymbol{z})$ and optimizing the evidence lower bound (ELBO):

$$\log p_\theta(x \mid \boldsymbol{c}) \;\geq\; \mathbb{E}_{q_\phi(\boldsymbol{z}|x,\boldsymbol{c})}[\log p_\theta(x \mid \boldsymbol{z}, \boldsymbol{c})] \;-\; \mathrm{KL}(q_\phi(\boldsymbol{z} \mid x, \boldsymbol{c}) \,\|\, p(\boldsymbol{z}))\,, \tag{2}$$

where q_ϕ denotes an approximate posterior (encoder) and $p_\theta(x \mid \boldsymbol{z}, \boldsymbol{c})$ the conditional likelihood (decoder). Introducing a multiplicative factor β on the KL term yields the β-VAE objective (Higgins et al., 2017).

Formally, the register of n_q qubits introduced in Section 1 is described by a unit-norm vector in the space $\mathcal{H} = (\mathbb{C}^2)^{\otimes n_q}$ of dimension $d = 2^{n_q}$, where $\otimes$ is the tensor (Kronecker) product. Such a single vector is a **pure state**; a **mixed state** is a probabilistic mixture of pure states, described by a positive semidefinite, unit-trace matrix ρ called the **density operator**. We write a state vector as $|\psi\rangle$, its conjugate transpose as $\langle\psi|$, and the computational basis vectors (the one-hot vectors of $\mathbb{C}^d$) as $|k\rangle$, $k = 0, \ldots, d-1$; $|0\rangle$ is the basis state with every qubit in 0. The previously described **amplitude encoding** (Section 2) embeds a classical vector in $\mathcal{H}$ by identifying an ℓ_2-normalized vector with the amplitude vector of a pure state, thereby representing d real-valued components using n_q qubits. The aforementioned parameterized quantum circuit (PQC), or **ansatz**, is formally a unitary matrix $U(\boldsymbol{\theta})$, i.e., $U^\dagger U = I$ with $U^\dagger$ the conjugate transpose (the **adjoint**), so that $U^\dagger$ undoes U; it is composed of alternating layers of trainable single-qubit rotations and fixed entangling operations. The trainable rotation mentioned in Section 2 is $R_Y(\theta) = \begin{pmatrix} \cos\frac{\theta}{2} & -\sin\frac{\theta}{2} \\ \sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix}$, a 2-D rotation of one qubit's amplitudes; the two-qubit gate CNOT flips a target qubit when its control qubit is 1, and the controlled rotation CRY of Equation 1 applies R_Y to a target only when its control is 1. Classical information is extracted through measurement, typically as expectation values $\langle O\rangle = \langle\psi| O |\psi\rangle$ of the Pauli observables behind the three measurement axes of Section 2, $X = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, $Y = \left(\begin{smallmatrix} 0 & -i \\ i & 0 \end{smallmatrix}\right)$, and $Z = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$, where $\langle Z_k\rangle$ denotes Z applied to qubit k. A single qubit's state can be drawn as the point $(\langle X\rangle, \langle Y\rangle, \langle Z\rangle)$ in the unit ball, the **Bloch sphere**: pure states lie on its surface, mixed states inside, and its center is the **maximally mixed** state. When qubits are entangled, each one taken alone is generally mixed; its **reduced** density operator is obtained by averaging out (partial-tracing) the other qubits.

A.2 Model

A.2.1 Molecular Representation

As described in Section 3.1, each molecule is represented as a SELFIES string (Krenn et al., 2020) and subsequently tokenized with respect to a fixed vocabulary comprising 25 tokens. The resulting sequence of integer token identifiers is embedded directly into an amplitude vector, which is zero-

padded to dimension $d = 2^{n_q}$ (with the padding symbol assigned to token 0, thereby yielding a natural padding convention) and ℓ_2-normalized:

$$|\psi_x\rangle \;=\; \sum_{k=0}^{d-1} x_k \,|k\rangle\,, \qquad \boldsymbol{x} \;=\; \boldsymbol{t}/\|\boldsymbol{t}\|_2, \tag{3}$$

The corresponding quantum state is prepared by a standard state-preparation circuit built from uniformly controlled rotations (Mottonen et al., 2004; Schuld & Killoran, 2019), which loads any given vector into the amplitudes. This encoding stage contains no trainable parameters. The use of raw token identifiers as amplitudes is an intentional simplification that fixes the number of classical pre-encoding parameters to zero.

A.2.2 Condition Injection

C-QOBRA is implemented without auxiliary condition qubits: each normalized condition variable $c_k \in [-1, 1]$ is incorporated via the controlled-rotation brickwork of Equation 1, applied sequentially as an even sublayer (CRY gates on qubit pairs $(0, 1), (2, 3), \dots$) followed by an odd sublayer (pairs $(1, 2), (3, 4), \dots$), the *brickwork* pattern of Appendix A.2.3, with all weights $w_{k,j}$ and biases $b_{k,j}$ being trainable parameters. Parameterizing the entanglement strength as an explicit function of the condition provides greater expressivity than architectures based on single-qubit condition rotations followed by fixed CNOT gates because the condition directly modulates inter-qubit correlations rather than only rotating individual qubits.

Condition injection is performed at $n_{\text{blocks}} + 1 + \text{cond_depth}$ locations per circuit. For the configuration considered here, this corresponds to four injection points in both the encoder and the decoder, each with independent parameterization. Tail injection points implement the adjoint layout $C^\dagger$ (the mirror-image gate order), obtained by reversing the sublayer order and the condition order, and negating the rotation angles; however, because the tail carries its own independently trained parameters, it does not simply invert C. Reintroducing the same condition information at each depth constitutes the aforementioned data re-uploading (Pérez-Salinas et al., 2020) and is essential for maintaining expressivity in $\boldsymbol{c}$ within a single-condition circuit. Finally, encoding the conditions within the data register eliminates two-qubit gates between the data qubits and separate condition qubits, keeps the circuit width fixed at n_q qubits, and reduces circuit depth. On real hardware, these properties limit *decoherence*, the loss of quantum information to the environment that grows with circuit depth, and reduce susceptibility to *barren plateaus*, the quantum analog of vanishing gradients, in which gradients shrink exponentially with the number of qubits (McClean et al., 2018).

A.2.3 Circuits

Encoder With $n_{\text{blocks}} = r$ blocks, the encoder and decoder unitaries are

$$U_{\text{enc}}(\boldsymbol{\theta}, \boldsymbol{c}) \;=\; C^\dagger(\boldsymbol{c})\, L(\boldsymbol{\theta}_{r+1}) \left[\prod_{\ell=1}^{r} C(\boldsymbol{c})\, E\, L(\boldsymbol{\theta}_\ell)\right] C(\boldsymbol{c}), \tag{4}$$

$$U_{\text{dec}}(\boldsymbol{\phi}, \boldsymbol{c}) \;=\; C^\dagger(\boldsymbol{c}) \left[\prod_{\ell=r}^{1} L(\boldsymbol{\phi}_\ell)\, E^\dagger\, C^\dagger(\boldsymbol{c})\right] L(\boldsymbol{\phi}_0)\, C(\boldsymbol{c}). \tag{5}$$

where $L(\boldsymbol{\theta}_\ell) = \bigotimes_{q=0}^{n_q-1} R_Y(\theta_{\ell,q})$ is a trainable single-qubit rotation layer (one R_Y per qubit, applied in parallel), E is the entangling unitary, and $C(\boldsymbol{c})$ is a condition-injection layer. For $n_{\text{blocks}} = r$, the encoder consists of an initial (head) condition injection, followed by r repetitions of the rotation then entanglement layer, a subsequent injection, and finally a terminal rotation layer followed by cond_depth adjoint (tail) injections. A single entangling block implements a complete brickwork pattern comprising an even sublayer acting on qubit pairs $(0, 1), (2, 3), (4, 5), (6, 7)$, followed by an odd sublayer acting on $(1, 2), (3, 4), (5, 6)$. Equivalently, it contains $n_q - 1 = 7$ parameter-free nearest-neighbor CNOT gates:

$$E \;=\; \underbrace{\prod_{q\in\{1,3,5\}} \text{CNOT}_{q,\,q+1}}_{\text{odd sublayer}} \;\; \underbrace{\prod_{q\in\{0,2,4,6\}} \text{CNOT}_{q,\,q+1}}_{\text{even sublayer}}. \tag{6}$$

This alternating rotation–entangler construction corresponds to the circuit-centric ansatz family of Schuld et al. (2020) and is hardware-efficient in the sense of Kandala et al. (2017): it uses two-qubit gates only between neighboring qubits on a line, the connectivity of many current devices, so no extra SWAP gates are needed to move qubit states between distant positions.

Decoder The decoder (Equation 5) mirrors the encoder: the circuit depth ordering is inverted, C is replaced by $C^\dagger$, and the brickwork connectivity is applied from odd-to-even rather than even-to-odd qubits. The decoder rotation angles are independent trainable parameters; i.e., they are not constrained to match those of the encoder. The ℓ_2-normalized vector of all 2^{n_q} output amplitudes forms the memory. We specify the two readout heads introduced in Section 3. The **register readout** splits the decoder's qubits into a group indexing sequence position and a group indexing the token and directly decodes the sequence from the resulting amplitudes; it introduces no additional classical parameters, requires $n_q = 10$, and is used in Section 4.2 and Figure 5(a). The **Transformer post-processor** is a single-layer `TransformerDecoder` (Vaswani et al., 2017) in the pre-layer-normalization arrangement of Xiong et al. (2020); it performs cross-attention over the memory, reshaped without parameters into $256/d_{\text{model}}$ vectors of width d_{model}, with no positional embeddings on either the memory or the target tokens. Its model width d_{model} is varied in Section 4.4, while the feed-forward sublayer dimensionality is fixed at $2d_{\text{model}}$. We also implemented a third head—namely, a tiled pure-quantum projection that factorizes across positions—but it is not used in the present experiments.

A.2.4 Latent Space

The latent register comprises the first $n_q - n_{\text{trash}}$ qubits, while the remaining n_{trash} qubits are the trash qubits described in Section 3.4 (Romero et al., 2017), trained to end in the fixed state $|0\rangle$ so that the input's information is compressed into the latent qubits. For $n_{\text{trash}} = 3$, the latent dimension is therefore 5. Since all circuit primitives (R_Y, CNOT, and CRY) are real and the amplitude-encoded input state is real, the state vector remains in the real subspace throughout the computation. Consequently, $\langle Y_i \rangle = 0$ identically, and each latent qubit is confined to the X–Z great circle of the Bloch sphere, implying a single real degree of freedom per qubit. We define the latent coordinate as

$$z_i \;=\; \operatorname{arctanh} \langle Z_i \rangle, \quad \langle Z_i \rangle = \langle \psi_x |\, U_{\text{enc}}^\dagger(\boldsymbol{\theta}, \boldsymbol{c})\, Z_i\, U_{\text{enc}}(\boldsymbol{\theta}, \boldsymbol{c})\, |\psi_x\rangle\,, \tag{7}$$

where $|\psi_x\rangle$ is the amplitude-encoded molecule. This chart corresponds to the inverse Gudermannian parametrization of the Bloch-sphere latitude (i.e., $z = \operatorname{arctanh}(\sin\lambda)$ for latitude λ), yielding a bijection $(-1, 1) \to \mathbb{R}$. This choice is motivated by two considerations. First, it provides a principled justification for an approximately Gaussian prior: by Archimedes' hat-box theorem, a Haar-uniform real qubit satisfies $\langle Z \rangle \sim \mathcal{U}(-1, 1)$, and the push-forward distribution under $\operatorname{arctanh}$ is $\text{Logistic}(0, \frac{1}{2})$, with a standard deviation $\pi/(2\sqrt{3}) \approx 0.907$, which is close to that of $\mathcal{N}(0, 1)$. Second, the mapping is exactly invertible for decoding: seeding the decoder with $\theta_i = \arccos(\tanh z_i) \in [0, \pi]$ enforces $\langle Z_i \rangle = \tanh z_i$, so that encoding and decoder seeding implement the same bijection in opposite directions. The trash qubits are not seeded and are instead initialized to $|0\rangle$.

This chart deteriorates when the latent qubits become mixed: as the latent register size increases, individual reduced states tend toward near-maximal mixing, $\langle Z \rangle$ concentrate near zero, $\operatorname{arctanh}$ become approximately linear, and posterior sampling noise can dominate the mean $\boldsymbol{\mu}$. To mitigate this effect, C-QOBRA introduces a purity regularizer $\sum_i (1 - |r_i|^2)$ over the latent qubits, where $|r_i|^2 = \langle X_i \rangle^2 + \langle Z_i \rangle^2$ denotes the squared Bloch radius. The variational posterior is then

$$q(\boldsymbol{z} \mid \boldsymbol{x}, \boldsymbol{c}) = \mathcal{N}\big(\boldsymbol{\mu},\, \operatorname{diag}(\boldsymbol{\sigma}^2)\big), \quad \mu_i = \operatorname{arctanh} \langle Z_i \rangle, \quad \log \boldsymbol{\sigma}^2 = f_\sigma(\boldsymbol{\mu}), \tag{8}$$

where f_σ is a compact $\tanh$-activated prediction head with a learned affine scaling. Latent codes are sampled using the reparameterization trick (Kingma & Welling, 2022), i.e., $\boldsymbol{z} = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\varepsilon}$ with $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$.

A.2.5 Conditional prior

During generation, latent samples $\boldsymbol{z}$ are drawn directly from the learned conditional prior $p(\boldsymbol{z} \mid \boldsymbol{c}) = \mathcal{N}(\boldsymbol{\mu}_p(\boldsymbol{c}), \boldsymbol{\sigma}_p^2(\boldsymbol{c}))$, and the encoder is not evaluated. By default, the mappings $\boldsymbol{\mu}_p$ and $\log \boldsymbol{\sigma}_p^2$ are parameterized by two compact feed-forward networks of the form $\text{Linear} \to \text{GELU} \to \text{Linear}$,

whose output layers are initialized to zero, thereby ensuring that optimization starts from the standard normal prior $\mathcal{N}(\mathbf{0}, \boldsymbol{I})$.

A.2.6 TRAINING OBJECTIVE

C-QOBRA is optimized in an end-to-end manner by minimizing the composite objective

$$\mathcal{L} \;=\; \mathcal{L}_{\text{recon}} \;+\; \beta\,\mathcal{L}_{\text{KL}} \;+\; \lambda_{\text{tr}}\,\mathcal{L}_{\text{trash}} \;+\; \lambda_{\text{pur}}\,\mathcal{L}_{\text{pur}} \;+\; \lambda_{\text{prop}}\,\mathcal{L}_{\text{prop}}, \tag{9}$$

where each component is summed over the mini-batch to ensure that the effective weighting of terms is invariant to batch size. The reconstruction loss is defined as the PAD-masked, token-level cross-entropy,

$$\mathcal{L}_{\text{recon}} = -\sum_{t=1}^{T} \log p(s_t \mid \boldsymbol{z}, \boldsymbol{c}, s_{<t}).$$

The latent regularization term corresponds to the closed-form KL divergence between Gaussians with respect to a learned conditional prior, with a per-dimension free-bits lower bound (Kingma et al., 2016):

$$\mathcal{L}_{\text{KL}} \;=\; \sum_{i} \max\big(\lambda_{\text{fb}},\, \text{KL}_i\big), \qquad \text{KL}_i = \text{KL}(q(z_i \mid \boldsymbol{x}, \boldsymbol{c}) \,\|\, p(z_i \mid \boldsymbol{c}))\,. \tag{10}$$

The trash loss follows the quantum autoencoder bottleneck proposed by Romero et al. (2017), and the purity loss implements the aforementioned Bloch-radius regularizer:

$$\mathcal{L}_{\text{trash}} \;=\; 1 - \langle 0|\,\rho_{\text{tr}}\,|0\rangle\,, \qquad \mathcal{L}_{\text{pur}} \;=\; \sum_{i} \big(1 - |r_i|^2\big), \tag{11}$$

with ρ_{tr} the reduced density operator of the three trash qubits, so that $\langle 0|\,\rho_{\text{tr}}\,|0\rangle$ is the probability of measuring all of them in 0. Finally, the optional property-matching term $\mathcal{L}_{\text{prop}} = (\mathbb{E}[\hat{y}] - y)^2$ is computed using a differentiable estimator implemented solely as a buffer and, therefore, introduces no additional trainable parameters.

A.3 PER-ARM GENERATION METRICS

Table 2 reports slope, MAPE, and reconstruction for every swept arm of all four descriptors. For Fsp3 we also report a hit rate: a sample counts as a hit if its Fsp3 lies within a *tolerance* of ± 0.05 of the target. Because this tolerance window overlaps the training distribution, hit rates are read against a condition-blind null, the hit rate of an unconditional sampler from the training distribution averaged over the six targets, which is 9.6% (Appendix A.7). Per-arm hit rates track the goodness-of-fit trends: on Fsp3 it rises from $13.4 \pm 4.3\%$ at 82 parameters to $32.8 \pm 2.8\%$ at 904 and $40.6 \pm 4.0\%$ at 3,344, then climbs more slowly to $54.8 \pm 6.2\%$ at the largest arm. We quote hit rates directly rather than as ratios to the condition-blind null, which is itself seed-dependent.

A.4 METRIC CHOICES

At $d_{\text{model}} = 2$ (82 parameters), one of the five Fsp3 seeds collapses: its medians remain constant at 0.50, yielding an estimated slope of 0.000; the corresponding R^2 is undefined because the response variable exhibits zero variance, leaving no variability for a linear model to attribute to the predictor. This seed is diagnostically useful because it reproduces the same failure mode as the unprocessed register readout, despite nominally including a post-processor, indicating that the observed transition depends on decoding-path capacity rather than on the mere inclusion of a classical module.

Although the generation logs report Pearson correlations between requested and achieved medians, we intentionally do not use this statistic in the analysis. When computed over six points, Pearson's r primarily rewards monotonic ordering rather than calibration or steering magnitude, and thus can overstate performance for ineffective arms. For example, the register readout—whose medians increase only from 0.50 to 0.60 across the full request range—achieves $r = 0.90$, while the $d_{\text{model}} = 4$ arm reaches $R^2 = 0.964 \pm 0.025$ despite operating at little more than half strength (slope 0.573 ± 0.104, MAPE $21.4 \pm 5.8\%$). Accordingly, we prioritize slope, MAPE, and achieved span as the discriminative metrics, as reported in Table 2.

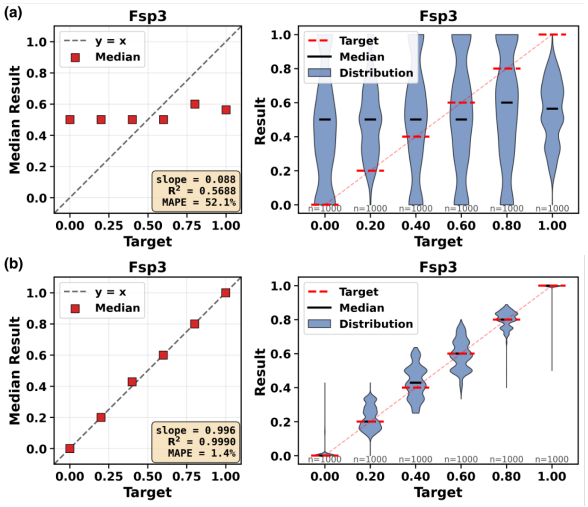

Figure 5: Fsp3-conditioned molecular generation (1,000 molecules per target). **(a)** Register-readout decoder without post-processing: the median predicted response remains approximately constant at ≈ 0.5 (slope 0.088; MAPE 52.1%), and the violin plots span nearly the entire $[0, 1]$ interval for all target values. **(b)** Identical quantum encoder augmented with a post-processing module of dimension $d_{\text{model}} = 128$: the median response closely follows the specified target (single seed shown; five-seed slope 0.963 ± 0.041, MAPE $4.2 \pm 2.2\%$, Table 2), and the distributions become tightly concentrated around the target value.

A.5 Descriptor vs. Register Readout

Comparing Figure 2(a) with the Fsp3 register-readout panel (Appendix Figure 5a) isolates the descriptor's effect: decoder, objective, and conditioning are identical, so differences reflect descriptor form rather than conditioning quality. Crippen logP is an additive sum of atom/fragment contributions (Wildman & Crippen, 1999) and is thus nearly linear in the token histogram: shifting marginals toward carbon- vs heteroatom-rich tokens changes logP monotonically without positional coordination. As a result, even crude marginal shifts yield ordered but under-scaled control (high R^2, low slope in Figure 2(a)).

Fsp3, by contrast, is a ratio (sp^3 carbons / total carbons) and is relatively invariant to size and composition; increasing it requires coordinated choices across multiple SELFIES positions (e.g., consistent aromatic vs non-aromatic tokens and matched ring openings/closures). Thus, logP is encoded in per-position marginals, whereas Fsp3 depends on inter-position correlations. Register readout decodes positions independently and drops these correlations, while a Transformer post-processor can exploit them via causal attention (e.g., linking ring closures to prior openings). Fsp3 is therefore present in the register state recoverable with a 904-parameter head (Section 4.4), but not accessible to an independent readout.

Implications: (1) post-processing requirements are property-dependent—additive descriptors remain controllable from marginals, while ratio-valued descriptors do not; (2) prediction—additive descriptors (e.g., molecular weight, heteroatom count) should degrade gracefully under register readout, whereas ratio/topology-valued ones (e.g., Fsp3, scaffold heteroatom fraction) should not. Appendix A.6 reports register-readout and swept arms for QED and fScafHetero alongside logP. fScafHetero is a second instance of the ratio/topology case, and a less composition-bound one, since it is defined on the scaffold and so depends on ring topology rather than on token identity alone: under the register readout its medians are flat (Figure 6), and they order only once a wide head is attached (Figure 3(d)), the same pattern Fsp3 shows. QED, also ratio-valued, stays weakly steered even at the widest head, so these arms are consistent with the prediction rather than a controlled test of it; the properties differ in more than functional form.

Table 2: Post-processor sweep for all four conditioning descriptors. d is the decoder width d_{model} and Params the post-decoder trainable parameter count. Entries are mean $\pm$ standard deviation over five seeds; rows marked † include collapsed runs (see Appendix A.6).

d	Params	Slope	MAPE (%)	Recon (%)
Fsp3				
2†	82	0.226 ± 0.215	43.9 ± 16.3	37.2 ± 1.5
4	260	0.573 ± 0.104	21.4 ± 5.8	46.3 ± 1.3
8	904	0.814 ± 0.014	10.5 ± 3.2	49.6 ± 0.3
16	3,344	0.847 ± 0.021	7.3 ± 2.5	51.8 ± 0.1
32	12,832	0.866 ± 0.025	5.6 ± 2.1	53.9 ± 0.2
64	50,240	0.894 ± 0.049	6.7 ± 1.7	55.5 ± 0.2
128	198,784	0.963 ± 0.041	4.2 ± 2.2	56.9 ± 0.3
QED				
2	82	0.025 ± 0.051	15.4 ± 0.8	36.3 ± 1.2
4	260	0.067 ± 0.035	14.9 ± 0.4	43.0 ± 1.3
8	904	0.182 ± 0.031	13.5 ± 0.6	47.4 ± 0.6
16	3,344	0.276 ± 0.072	12.2 ± 1.0	49.6 ± 0.1
32	12,832	0.295 ± 0.021	12.1 ± 0.2	51.5 ± 0.2
64	50,240	0.318 ± 0.025	12.0 ± 0.5	53.5 ± 0.2
128	198,784	0.395 ± 0.042	11.1 ± 0.8	55.2 ± 0.2
logP				
2	82	0.136 ± 0.155	88.0 ± 16.4	37.0 ± 1.1
4	260	0.618 ± 0.032	42.3 ± 3.3	44.7 ± 1.3
8	904	0.742 ± 0.015	32.0 ± 1.3	48.8 ± 0.2
16	3,344	0.788 ± 0.011	29.5 ± 0.5	50.8 ± 0.2
32	12,832	0.820 ± 0.019	28.4 ± 0.9	52.4 ± 0.2
64	50,240	0.848 ± 0.007	26.2 ± 0.7	54.0 ± 0.2
128	198,784	0.879 ± 0.028	25.6 ± 0.7	55.5 ± 0.1
fScafHetero				
2†	82	0.054 ± 0.121	95.8 ± 9.3	36.5 ± 2.2
4	260	0.327 ± 0.196	61.2 ± 22.3	44.2 ± 1.6
8	904	0.595 ± 0.139	52.4 ± 20.3	47.7 ± 0.3
16	3,344	0.908 ± 0.059	29.8 ± 14.5	50.2 ± 0.2
32	12,832	0.840 ± 0.097	16.0 ± 4.6	52.1 ± 0.3
64	50,240	0.777 ± 0.342	18.5 ± 12.1	53.9 ± 0.2
128	198,784	0.950 ± 0.079	14.9 ± 8.9	55.6 ± 0.2

A.6 Additional conditioning results

This appendix reports post-processor sweeps for logP, QED, and fScafHetero alongside Fsp3 (all in Table 2). Each arm used five seeds and one decoder layer, generating $1{,}000$ molecules per target; results are the mean $\pm$ standard deviation across seeds. Because the seed also sets the train/test split, the spread reflects both split and initialization variability. Slope and MAPE are computed by fitting the median achieved value to the requested target, and Recon is token-level reconstruction accuracy on held-out molecules. Rows marked † include at least one collapsed seed (constant output): 1/5 for Fsp3 and 4/5 for fScafHetero at $d_{\text{model}} = 2$. For collapsed seeds, we report a slope of 0. Target grids differ by property: Fsp3 $\{0.0, 0.2, \ldots, 1.0\}$, logP $\{-3, -2, -1, 0, 1, 2, 3\}$, QED $\{0.30, 0.35, 0.40, 0.45, 0.50\}$, and fScafHetero $\{0.0, 0.2, 0.3, 0.4, 0.5\}$. Since MAPE is relative and target magnitudes differ (and logP crosses zero), compare MAPE only within a property. Register-readout steering arms reported here remain single-seed.

A.6.1 logP

Table 2 gives the logP sweep. It extends the two-readout comparison of Figure 2 across the size axis: the register-readout arm is already well ordered but under-scaled, and as the post-decoder widens, the slope rises, MAPE falls, and the achieved range widens on both sides of the training center. The sweep flattens at the top, as it does on Fsp3.

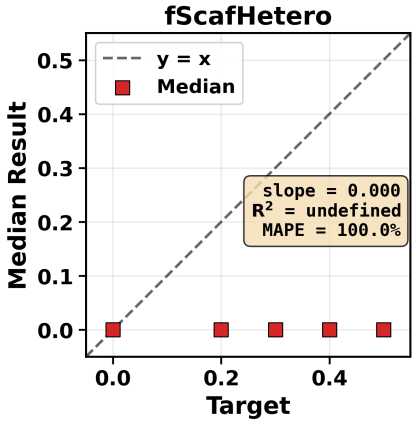

Figure 6: fScafHetero under the register readout, with no post-processor: median achieved value against the requested target, with the fit printed inside the panel. The medians are constant, which is the degenerate fit noted in Table 2; the corresponding widest swept arm is Figure 3(d). The panel shows medians only and says nothing about spread.

A.6.2 QED

Steering is weak in every QED arm, and we include Table 2 as evidence that the property is hard under this model rather than as a demonstration that the method transfers to it. The largest slope anywhere in the table is 0.395 ± 0.042, and the achieved means stay within roughly 0.37-0.45 against a request that spans 0.30-0.50; no arm should be described as steering QED well. The QED arms were regenerated on the 0.30-0.50 grid quoted above, which is the grid reported here. Figure 3(b) shows the widest arm: the medians move in the requested direction but traverse only a fraction of the requested interval.

A.6.3 fScafHetero

Two features of this run matter for reading Table 2. First, the 0.1 target was dropped before generation because the training set had no support below it within the tolerance window, so five targets were scored instead of six. Second, the register-readout arm has a degenerate fit: its medians are constant across all five targets, so the slope is 0.000 and R^2 is undefined rather than zero, and the table marks it with a dagger. With a post-processor in the decode path, the fit becomes ordered, and the achieved range opens up, as Figure 6 and Figure 3(d) show for the two ends of the sweep.

Read together with the main text, these three sweeps are consistent with the account in Section 4.3: the additive descriptor is the one that the parameter-free readout orders correctly on its own, while the two ratio-valued descriptors are the ones on which it does not; in every case, the post-processor is what converts a register state into a steerable output. They are not a controlled test of that account since the properties differ in more than just their functional forms.

A.7 Training Distribution & Condition-Blind Null

Figure 7 presents the training-set distribution of Fsp3 on QM9, thereby specifying the condition-blind null described in Appendix A.3. The per-target training support within a tolerance window of ± 0.05 equals 13.49, 14.08, 10.14, 8.55, 5.53, and 5.68% for targets spanning 0.0 to 1.0, respectively; the null level is defined as their arithmetic mean, 9.6%.

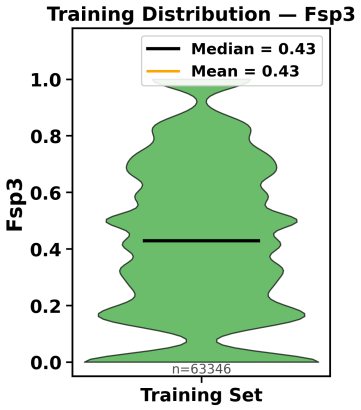

Figure 7: Training-set distribution of Fsp3 on QM9.