

A trainability-neutral occupation-entropy regularizer against mode collapse in quantum generative models

Jian Feng Kong,¹ Siong Thye Goh,¹ Chee Kwan Gan,¹ Stefano Carrazza,^{2,3,4} Jun Yong Khoo¹

¹Institute of Advanced Intelligence and Computing (IAIC), Agency for Science, Technology and Research (A*STAR), 1 Fusionopolis Way, #16-16 Connexis, Singapore 138632, Republic of Singapore

²Dipartimento di Fisica, Università degli Studi di Milano, Milan, Italy; ³INFN, Sezione di Milano, I-20133 Milan, Italy

⁴Quantum Research Center, Technology Innovation Institute, Abu Dhabi, UAE

July 8, 2026

Abstract

Quantum generative models face two well-known obstacles: barren plateau which affects the trainability of the model; and mode collapse is the failure of a trained model to reproduce all of the target distribution’s modes. The two are sometimes linked, because the low-resolution cost functions used to avoid barren plateaus are also the ones least able to tell which modes carry probability. In this work, we introduce a target-free coverage regularizer: a penalty against mode collapse that uses only the model’s own probability in each mode and the fixed geometry of the modes, and never the target mode populations one is trying to learn. Its central property is that it is *trainability-neutral*: because the observables that read off the mode occupations act on only a few qubits (they have low Pauli weight), adding the regularizer does not worsen the barren plateau. On a molecular-conformation benchmark whose modes are physically defined and exactly countable, it recovers about 80% of the modes lost by a collapsing quantum generative adversarial network (QGAN) at no trainability cost. The method also has a definite limit, which we characterize: how well a starved mode can be refilled is set by how strongly the base cost concentrates probability away from it (a quantity labelled g_{base}); recovery is effective when g_{base} is small but is exponentially suppressed when g_{base} is large. Energy-based generation, whose g_{base} is more than an order of magnitude larger than the adversarial case, is therefore beyond the regularizer’s reach.

1 Background

A quantum generative model prepares a parametrized state $|\psi(\boldsymbol{\theta})\rangle$ and samples it; the measurement-outcome probabilities $p_{\boldsymbol{\theta}}(x) = |\langle x|\psi(\boldsymbol{\theta})\rangle|^2$ define the model distribution (a circuit Born machine). Two well-known obstacles stand between such a model and a useful result. The first is the barren plateau: for expressive circuits the gradient of a typical cost concentrates exponentially in the qubit number, and this concentration is cost-function dependent, so global or high-resolution costs plateau while local, low-resolution costs stay trainable [1–3]. The second is mode collapse, the signature failure of quantum generative adversarial networks and of recent quantum molecular generators [4–6]. The two interact: a cost kept local enough to avoid the plateau is, for that same reason, blind to which modes hold probability, so a fix for one pathology can potentially worsen the other.

In this work, we target molecular conformational ensembles as an illustrative example. Here a mode is a physically defined metastable state (a basin of a torsional free-energy landscape) with a known relative population, so “did the model cover the modes?” is an exact count. Previous works such as QCA-MolGAN [5] and the Latent Style-based Quantum Wasserstein GAN [6], reduce collapse in molecular-graph space, but neither characterizes the relationship between barren plateaus and mode collapse nor provides a trainability guarantee. Those two gaps are what we fill here.

2 Methodologies

The system is a discretized torsional energy landscape: a grid coordinate x labels either one torsion angle ϕ (synthetic multi-well targets and butane) or the Ramachandran pair (ϕ, ψ) (alanine dipeptide), and the circuit Born machine assigns probabilities $p_{\boldsymbol{\theta}}(x)$ to these cells. The chemistry enters through the




Figure 1: *Left*: where barren-plateau-free quantum cost functions lose modes ($B=16$, heavy-tailed). MMD is faithful; a local-marginal cost drops minor modes; a capacity-limited QGAN collapses; energy-based generation mode-seeks. *Right*: the occupation-entropy regularizer recovers most of the QGAN collapse ($\sim 80\%$) at no trainability cost.

potential $V(x)$, whose local minima define metastable conformational basins (for example, *trans/gauche* in butane and the standard Ramachandran basins in alanine); these basins are fixed before training and constitute the mode geometry. Group the grid cells into these B modes, and let $m_b = \sum_{x \in b} p_{\theta}(x)$ be the model’s occupation of mode b (the probability it assigns to that basin). A regularizer is *target-free* if it is a function of the occupation vector $\{m_b\}$ and the mode geometry alone, and not of the target populations q_b . We study two regularizers. The occupation-entropy regularizer $R_S = -\lambda S(\{m_b\})$, with $S(\{m_b\}) = -\sum_b m_b \log m_b$ the Shannon entropy, rewards spreading probability across the geometrically defined modes. The minimum-occupation penalty $R_{\tau} = \lambda \sum_b [\max(0, \tau - m_b)]^2$ activates only when a mode’s occupation falls below a threshold τ . Both depend on p_{θ} only through the $B = \mathcal{O}(1)$ occupations, so both are cheap and, as we show next, harmless to trainability.

Trainability-neutrality. Each occupation $m_b = \langle O_b \rangle$ is the expectation of an indicator O_b that acts on only a few high-order qubits of one torsion, a low-Pauli-weight observable, which by the cost-function-dependent barren-plateau result [2] is resistant to the plateau. Its gradient variance therefore decays no faster than the base cost’s, and a triangle inequality on the two jointly distributed gradients gives, when the regularizer gradient is subdominant at initialization, an equal decay exponent for the base cost and the regularized cost: adding the regularizer does not steepen the gradient-variance decay. We confirm this numerically, including in the adversarial regime; where the regularizer is not strictly subdominant it flattens the plateau (its low-weight gradient decays more slowly than a low-gradient base cost), so it is at worst neutral and sometimes beneficial. This is the central property of the method: mode coverage can be enforced without paying a trainability cost.

The regularizer does not repair every collapse: because the occupations lie on the probability simplex $\sum_b m_b = 1$, the stationary point of the entropy-regularized cost is a Gibbs (softmax) distribution over modes with “energies” $g_b = \partial L_{\text{base}} / \partial m_b$ and a temperature set by λ , so a starved mode settles to

$$m^* = C e^{-g_{\text{base}}/\lambda}, \quad C = 1/Z, \quad (1)$$

where g_{base} measures how strongly the base cost concentrates probability away from that mode and Z normalizes over modes. Coverage is restored when g_{base} is small ($g_{\text{base}} \ll \lambda$) and is exponentially suppressed when g_{base} is large: refilling a mode that the base cost drives hard toward zero would require a λ so large it distorts the rest of the distribution.

3 Results and Discussion

The benchmark uses synthetic multi-well torsion targets (mode count $B \in \{8, 12, 16\}$; either balanced, or heavy-tailed, meaning a few dominant modes plus several minor ones) with a genuine potential, plus butane ($B=3$) and alanine dipeptide ($B=5$) as molecular anchors. It first locates where barren-plateau-free costs lose modes (Fig. 1, left). Maximum mean discrepancy (MMD) with a characteristic




Figure 2: The regularizer’s coverage limit. The occupation of the most-starved mode after training follows the Gibbs form $m^* = (1/Z) e^{-g_{\text{base}}/\lambda}$, so $-\log m^*$ is linear in $1/\lambda$ with slope g_{base} ; the slope is measured independently of the fit. For the capacity-limited QGAN $g_{\text{base}} \approx 0.07$ (nearly flat), so coverage is recovered; for energy-based generation ($\beta \rightarrow 0$) $g_{\text{base}} \approx 3.2$ (more than an order of magnitude larger), beyond the regularizer’s reach. Open triangles (energy, at $1/\lambda=5, 10$) mark near-zero occupations ($m^* < 10^{-9}$) too small to resolve reliably, excluded from the fit.

kernel (one whose discrepancy is zero only at $p_\theta=q$) is robustly mode-faithful and its minimum is $p_\theta=q$ for any bandwidth. A local single-qubit-marginal cost drops minor modes. Energy-based generation, which minimizes a variational free energy $F = \langle V \rangle - TS$ estimated from the model, mode-seeks: at high effective temperature ($\beta \rightarrow 0$) it concentrates onto the single deepest well (reverse-Kullback–Leibler collapse). A QGAN, in the standard mode-collapse regime of a capacity-limited generator outpaced by its discriminator, collapses robustly: collapse rate 1.0 (Wilson 95% confidence interval $[0.84, 1.0]$, 20 seeds), with the lost-mode count growing from 3.5/8 to 7.7/16, and cost-induced, since a same-capacity MMD loses at most 1.7 modes.

The occupation-entropy regularizer recovers most of the QGAN collapse (Fig. 1, right; $7.7 \rightarrow 1.2$ dropped modes at $B=16$) at no trainability cost. This is the regime the method targets: the adversarial cost concentrates only weakly on the starved mode, so its g_{base} is small. We note that while coverage is recovered, there is a mild accuracy cost (a slight over-spreading of probability) and a residual collapse rate, so this is a substantial mitigation rather than a complete cure. A complementary minimum-occupation penalty is the accuracy-preserving choice for gentle minor-mode dropout: on the local-marginal cost it restores coverage with less over-spreading than occupation-entropy ($L_1 \approx 0.42$ versus 0.62), but it cannot overcome the strong adversarial concentration of the QGAN.

To elucidate the performance limitation of this method, we measure g_{base} independently of any fit, as a mass-conserving directional derivative of the base cost in occupation space, and compare it to the slope of $-\log m^*$ versus $1/\lambda$ (Fig. 2). For energy-based generation the two agree quantitatively (fitted slope 3.22; independent measurement near collapse ≈ 3.5 ; analytic well-depth gap $\Delta V/k_B T = 3.04$); for the QGAN both are small and agree to order of magnitude (fitted 0.07, independent 0.12 ± 0.07). The separation $g_{\text{base}}^{\text{QGAN}} \approx 0.1 \ll g_{\text{base}}^{\text{energy}} \approx 3$ (about 30 $\times$) implies that the QGAN’s collapse is recovered at a modest λ , whereas energy-based mode-seeking, driven by a deep well-depth gap, is beyond the regularizer’s reach at any practical λ .

In summary, we introduced a target-free, trainability-neutral coverage regularizer against mode collapse in quantum generative models and evaluated it on a conformational benchmark whose modes are physically defined and exactly countable. We mapped where barren-plateau-free cost functions lose modes, showed that an occupation-entropy regularizer recovers most of the adversarial (QGAN) collapse at no trainability cost, and characterized when such a regularizer can and cannot help through the base cost’s concentration gradient g_{base} .

References

- [1] 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, 2018. doi:[10.1038/s41467-018-07090-4](https://doi.org/10.1038/s41467-018-07090-4). [arXiv:1803.11173](https://arxiv.org/abs/1803.11173).
- [2] M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature Communications*, 12(1):1791, 2021. doi:[10.1038/s41467-021-21728-w](https://doi.org/10.1038/s41467-021-21728-w). [arXiv:2001.00550](https://arxiv.org/abs/2001.00550).
- [3] Manuel S. Rudolph, Sacha Lerch, Supanut Thanasilp, Oriel Kiss, Oxana Shaya, Sofia Vallecorsa, Michele Grossi, and Zoë Holmes. Trainability barriers and opportunities in quantum generative modeling. *npj Quantum Information*, 10:116, 2024. doi:[10.1038/s41534-024-00902-0](https://doi.org/10.1038/s41534-024-00902-0). [arXiv:2305.02881](https://arxiv.org/abs/2305.02881).
- [4] Aaron Mark Thomas, Harry Youel, and Sharu Theresa Jose. VAE-QWGAN: Addressing mode collapse in quantum GANs via autoencoding priors. *Quantum Machine Intelligence*, 2025. doi:[10.1007/s42484-025-00314-z](https://doi.org/10.1007/s42484-025-00314-z). [arXiv:2409.10339](https://arxiv.org/abs/2409.10339).
- [5] Aaron Mark Thomas, Yu-Cheng Chen, Hubert Okadome Valencia, Sharu Theresa Jose, and Ronin Wu. QCA-MolGAN: Quantum circuit associative molecular GAN with multi-agent reinforcement learning. *Proceedings of IEEE Quantum Artificial Intelligence*, 2025. [arXiv:2509.05051](https://arxiv.org/abs/2509.05051).
- [6] Julien Baglio, Yacine Haddad, and Richard Polifka. Latent style-based quantum Wasserstein GAN for drug design. *arXiv preprint*, 2026. [arXiv:2603.22399](https://arxiv.org/abs/2603.22399).