

A wider net for new molecules.

ORCA's quantum prior makes generative chemistry models more diverse and more efficient, expanding the pool of viable discovery candidates.

QUANTUM AI FOR SCIENCE

+49%

more valid & unique molecules vs classical baseline on ORCA's PT-2 photonic quantum accelerator

+41%

more novel molecules not in the training data, directly expanding the early-stage discovery candidate pool

+23%

improvement in chemical authenticity, measured by Fréchet Chemical Distance (FCD)

1. Limited molecular diversity

Classical generative models cluster around known chemical space, missing the structural diversity that leads to novel hits.

Simple, factorisable distributions limit how much variety a finite-capacity model can produce, especially for data governed by quantum mechanics. The result is fewer valid, novel candidates and weaker coverage of chemical space.

2. The value that ORCA delivers

ORCA replaces the classical prior with a quantum prior from the PT-2 photonic quantum accelerator, exploring more of the chemical space that matters in every generation cycle. It works alongside your current workflow: same models, same training, same Python frameworks. Only the diversity and quality of what comes out changes.

More diverse, more authentic molecules per cycle mean a larger pool of viable discovery candidates.

3. Benchmark results

On the QM9 quantum chemistry dataset, a quantum prior outperformed every classical baseline tested, including Bernoulli and distinguishable-photon priors.

[Read the paper →](#)

USE-CASE IMPLEMENTATION

Hybrid solution

PT-2 photonic quantum accelerator (48 qumodes)

Classical HPC with NVIDIA GPUs

ORCA SDK with Python and PyTorch compatibility

AI approach

Quantum-accelerated GAN for molecular generation

Benchmarked on the QM9 quantum chemistry dataset