

Generative Adversarial Networks for Resource State Generation

Shahbaz Shaik,^{1,2,*} Sourav Chatterjee^{3,†}, Sayantan Pramanik,^{3,‡} and Indranil Chakrabarty^{1,2,§}

¹Centre for Quantum Science and Technology, International Institute of Information Technology, Hyderabad

²Center for Security Theory and Algorithmic Research,
International Institute of Information Technology, Hyderabad

³TCS Research, TATA Consultancy Services, India

We introduce a physics-informed Generative Adversarial Network framework that recasts quantum resource-state generation as an inverse-design task. By embedding task-specific utility functions into training, the model learns to generate valid two-qubit states optimized for teleportation and entanglement broadcasting. Comparing decomposition-based and direct-generation architectures reveals that structural enforcement of Hermiticity, trace-one, and positivity yields higher fidelity and training stability than loss-only approaches. The framework reproduces theoretical resource boundaries for Werner-like and Bell-diagonal states with fidelities exceeding 98%, establishing adversarial learning as a lightweight yet effective method for constraint-driven quantum-state discovery. This approach provides a scalable foundation for automated design of tailored quantum resources for information-processing applications, herein exemplified with teleportation and broadcasting of entanglement, and also opens up the possibility of using them in efficient quantum network design.

I. INTRODUCTION

Resource generation is a central challenge across the physical sciences. In classical contexts, it appears as the design of systems with prescribed responses—from photonic crystals that mold the flow of light [1] to metamaterials exhibiting negative refractive index [2]. Traditional parametric and variational methods become rapidly intractable as the number of degrees of freedom grows [3]. **Machine Making** (MK) has therefore emerged as a natural paradigm for inverse design, with neural networks (NNs) enabling automated discovery of structures in optics [4], acoustics [5], and materials science [6]. More generally, NNs are being widely employed across various areas of physics [6–11] for classification tasks [12], where the objective is to assign labels to data samples, and for generative modeling [13, 14], where new data are produced after learning underlying data distributions [15]. Generative Adversarial Networks (GANs) [16] offer an appealing framework for the latter purpose.

In quantum information science, many general challenges in characterizing and controlling quantum systems can be reformulated as machine-learning tasks. For instance, NN-based methods have been successfully applied to address foundational questions [17], phase transition detection [18], entanglement [19, 20] and nonclassicality [21, 22] identification, quantum experiment design [23], error correction [24, 25], and device calibration [26]. In particular, the inverse-design viewpoint (mentioned above) applies to quantum states themselves, which act as operational resources for information-processing tasks [27]. A valid

quantum state must satisfy Hermiticity, unit trace, and positive semidefiniteness [28], and its resourcefulness is often defined by protocol-dependent criteria such as teleportation fidelity [29] or broadcasting capability [30]. The analytical characterization of all states fulfilling a given criterion is rarely feasible, especially beyond the two-qubit regime [31]. This makes the systematic discovery of useful resource states a natural candidate for data-driven approaches that can explore and learn such constrained manifolds directly [15].

Quantum information processing often **does not** rely critically on entanglement as a resource. However, this resource can at times be task specific [32, 33] and sometimes beyond states and channel, i.e., in network where it can be topology specific [34]. It can be activated under global unitary operation [35] and can be conserved or destroyed under the action of channels [36, 37]. Maximally entangled Bell states enable perfect teleportation [29, 38]. However, there can be other entangled states that can act as a resource in offering the quantum advantage in teleportation (i.e., surpassing its classical limits) [29]. On the other hand, broadcasting of quantum correlation is a method to quantify the range of input parameters of a shared state under which it can act as a resource by broadcasting its quantum correlation into multiple weaker correlated states [30]. However, a more general resource theoretic framework was lately used to understand different types of resources including distributed quantum computation [39], superdense coding [40], distillation [41] and many more [33, 42]. The ability to identify and generate these resource states is hence essential.

Originally developed to learn complex data distributions, GANs have since been applied across physics, from phase-transition discovery [43] to particle-collision modeling [44]. Quantum variants (qGANs) [45, 46] and classical-quantum hybrids [15, 47] have demonstrated their ability

* shahbaz.shaik@research.iit.ac.in

† sourav.chat@tcs.com

‡ sayantan.pramanik@tcs.com

§ indranil.chakrabarty@iiit.ac.in

to reproduce valid quantum states, yet targeted generation of resource states optimized for specific tasks remains largely unexplored. Compared with likelihood-based or diffusion models [48], adversarial learning provides a lightweight mechanism in which the discriminator naturally enforces physical validity, making GANs particularly suited for low-dimensional, constraint-dominated problems such as synthesis of density matrices representing quantum states [28].

In summary, resource-state generation is a unifying inverse-design problem across physics, from creating metamaterials [] with target electromagnetic responses [3, 49] to engineering molecular structures for catalysis []. In quantum information science, this challenge is amplified: entangled states enable core protocols such as teleportation and broadcasting, yet analytical characterization of useful non-maximally entangled states remains elusive [27]. In this work, we present a GAN framework that treats quantum resource-state generation as an inverse-design problem. By embedding protocol-specific utility functions directly into training, our model learns distributions of valid two-qubit states acting as a resource for teleportation and broadcasting. We compare three generator architectures—two enforcing physical constraints through decomposition and one unconstrained model trained via loss penalties—and show that structural enforcement yields higher fidelity and training stability. Our results demonstrate adversarial learning as a practical and extensible method for generating tailored quantum resources. Generated quantum resources advantageous for communication can be used for large scale quantum network design [34, 50].

Our manuscript is organised as follows. In Section II, we provide a brief overview of the process of broadcasting entanglement and executing teleportation. Therein, we also define the resource states involved in these processes. In Section III, we describe the structure of GAN and the evaluation metrics associated with it. In Section IV, we generate resource states through GAN and compare them with the existing theoretical results. Thereafter, in Section V, we provide the concluding remarks, and finally in Section VI, we discuss the future research efforts that can be further explored in this direction.

II. QUANTUM RESOURCE STATES

Quantum resource states underpin numerous tasks in quantum information processing. In this section, we outline two pivotal protocols—broadcasting of quantum correlations and quantum teleportation—and clarify the criteria that render a quantum state a useful resource for each. The identification and synthesis of such states is central to the GAN-based framework developed in this work.

A. Resource States for Broadcasting

Broadcasting of quantum correlations is designed to increase the availability of entanglement across a network, effectively decomposing entanglement to generate more entangled pairs from an initial quantum resource [30, 51, 52]. Although typically demonstrated for two-qubit systems, the concept generalizes to higher-dimensional and multipartite settings, and has been extended to resources beyond entanglement, such as quantum discord and coherence [53–55]. Quantum operations for broadcasting are **unusually** implemented via cloning but not always restricted to cloning [56]. In this we consider broadcasting with the help of both local and nonlocal cloning.

a. Local Cloning. Consider an entangled state ρ_{12} shared between Alice and Bob. Local broadcasting introduces ancilla qubits $|\Sigma_3\rangle, |\Sigma_4\rangle$, followed by local cloning operations (U_1 at Alice’s site and U_2 at Bob’s). The output state is constructed as shown in Eq. (1):

$$\rho'_{1234} = (U_1 \otimes U_2)(\rho_{12} \otimes |\Sigma_3\rangle \otimes |\Sigma_4\rangle). \quad (1)$$

Tracing out the corresponding qubits, one obtains the local marginals in Eq. (2) and nonlocal marginals in Eq. (3):

$$\begin{aligned} \rho'_{13} &= \text{Tr}_{24}(\rho'_{1234}), & \rho'_{24} &= \text{Tr}_{13}(\rho'_{1234}), & (2) \\ \rho'_{14} &= \text{Tr}_{23}(\rho'_{1234}), & \rho'_{23} &= \text{Tr}_{14}(\rho'_{1234}). & (3) \end{aligned}$$

Entanglement is considered *broadcast non-optimally* if the nonlocal pairs ρ'_{14}, ρ'_{23} remain entangled [51]. For *optimal broadcasting*, it is additionally required that the local pairs ρ'_{13}, ρ'_{24} are separable [30, 52].

b. Nonlocal Cloning. When both qubits are co-located (are in the same lab), a global unitary U_{12} can be applied, as shown in Eq. (4):

$$\rho'_{1234} = U_{12}(\rho_{12} \otimes |\Sigma_3\rangle \otimes |\Sigma_4\rangle). \quad (4)$$

Local and nonlocal marginals are similarly defined by partial traces. The same entanglement criteria apply for optimal vs. non-optimal broadcasting.

c. Entanglement Detection. In 2×2 and 2×3 systems, the Peres-Horodecki criterion provides a necessary and sufficient test: if the partial transpose of ρ yields a negative eigenvalue, the state is entangled [57, 58]. Here we use this criteria to detect whether the output states are entangled or not.

Broadcasting has also been generalized to resources beyond entanglement—such as NPT states, coherence, and Bell nonlocality—and performed using asymmetric cloning, higher-dimensional operations, and continuous-variable systems [53, 54, 56, 59, 60]. Not every mixed state is broadcastable [61]; quantum mechanics fundamentally limits perfect cloning and broadcasting as compared to classical information theory [61, 62].

B. Resource States for Teleportation

Quantum teleportation exploits entanglement to transfer quantum information between parties. If Alice and Bob share a maximally entangled Bell state, such as the one in Eq. (5),

$$|\Psi_+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle), \quad (5)$$

then an unknown qubit can be teleported with fidelity $F = 1$ [29, 38]. More generally, the teleportation fidelity depends on the entanglement quality of the shared resource. For example, in the absence of entanglement, the maximal achievable fidelity is $F = 2/3$ [29]. Thus, any state giving $F > 2/3$ constitutes a resource for quantum teleportation.

Given a two-qubit density matrix of the general form in Eq. (6),

$$\rho = \frac{1}{4} \left[I_4 + \sum_i \sigma_i \otimes I + \sum_j I \otimes \sigma_j + \sum_{i,j} t_{ij} \sigma_i \otimes \sigma_j \right], \quad (6)$$

(where σ_i are Pauli matrices and $T = [t_{ij}]$ is the correlation matrix), the teleportation criterion is given by Eq. (7):

$$F_{\max}(\rho) = \frac{1}{2} \left[1 + \frac{1}{3} N(\rho) \right], \quad N(\rho) = \text{Tr} \left(\sqrt{T^\dagger T} \right) \quad (7)$$

States with $N(\rho) > 1$ or equivalently $F_{\max}(\rho) > \frac{2}{3}$ provide quantum advantage [29, 63–65]. Note that not every entangled state is useful for teleportation; counterexamples are documented [63, 64, 66]. The idea of a resource state in context is also being extended for a quantum network, by quantifying networks ability to do distributed teleportation [34, 67].

III. OVERVIEW OF GAN

A GAN consists of two neural networks in competition: a generator $G(z)$ that maps latent noise z to candidate states, and a discriminator $D(x)$ that distinguishes generated states from training samples [16, 68]. Training is framed as the minimax game shown in Eq. (8):

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim P_{\text{data}}} [\log D(x)] + \mathbb{E}_{z \sim P_z} [\log(1 - D(G(z)))]. \quad (8)$$

Intuitively, training proceeds as a two-player game: the discriminator improves at distinguishing training from generated samples, while the generator adapts to produce outputs that the discriminator cannot reject. At equilibrium, the generator’s output distribution matches the training distribution, enabling synthesis of new samples

with the same statistical properties. For density-matrix generation, this dynamic is particularly advantageous: the discriminator implicitly learns the manifold of valid quantum states, providing a learned enforcement of physicality that complements explicit constraint terms in the loss function.

A schematic overview of our physics-informed training pipeline is shown in Figure 1. During training, the generator maps latent samples z to candidate density matrices $\rho_{\text{gen}} = G(z)$. The discriminator compares ρ_{gen} with training samples ρ_{train} to produce the adversarial signal, while a separate physics module computes differentiable penalties for trace-one normalization, positive semidefiniteness, and task utility (broadcasting or teleportation). The generator is updated using the resulting weighted objective (Eq. (9)). After training, the generator is used as a sampler: drawing z yields $\rho_{\text{resource}} = G^*(z)$, a candidate resource state.

A. Conditional GANs and Quantum Constraints

For quantum state generation, validity requires Hermiticity, trace-one normalization, and positive semidefiniteness. A conditional GAN (CGAN) [69] extends the standard GAN by incorporating auxiliary information into both generator and discriminator, enabling targeted generation conditioned on external signals. In our setting, these signals take the form of physics-informed losses: trace loss, positive semi-definite (PSD) loss, and utility-specific losses (broadcasting vs. teleportation), which enforce both physical validity and task relevance. Figure 1 provides a schematic representation of the physics-informed CGAN framework used throughout this work.

1. Generator Architectures

We compare three generator designs that differ in how they enforce physicality. Models 1–2 impose Hermiticity and positive semidefiniteness by construction via matrix factorizations, and enforce the trace-one constraint via normalization. Model 3 outputs a Hermitian matrix directly and relies on loss penalties to learn physical validity.

- *Cholesky Decomposition* (Model 1): the network outputs the entries of a lower-triangular factor L (implemented by filling the lower triangle from a real-valued parameter vector). We assemble $\tilde{\rho} = LL^\dagger$, which is Hermitian and PSD by construction, and enforce trace-one by $\rho = \tilde{\rho}/\text{Tr}(\tilde{\rho})$.
- *Enhanced LDL Decomposition* (Model 2): the network outputs a unit-lower-triangular L and a diagonal D with strictly positive entries (implemented with a small ϵ floor for numerical stability), and constructs $\tilde{\rho} = LDL^\dagger$, followed by trace normalization.

- *Direct Matrix Generator* (Model 3): the network outputs an unconstrained real vector that is reshaped into the real and imaginary parts of a 4×4 matrix, which is then symmetrized to obtain a Hermitian $\tilde{\rho}$. Trace-one and PSD are not guaranteed by construction and are instead encouraged purely through the explicit loss terms.

Full architectural specifications and training hyperparameters are provided in Appendix A.

2. Adversarial Training

The generator and discriminator are trained in competition. The discriminator is a multilayer perceptron (MLP) with Leaky Rectified Linear Unit (ReLU) activations and dropout regularization. The generator is optimized using the composite loss in Eq. (9):

$$G_{\text{loss}} = L_{\text{adv}} + \lambda_1 L_{\text{trace}} + \lambda_2 L_{\text{psd}} + \lambda_3 L_{\text{task}}, \quad (9)$$

where L_{task} represents teleportation fidelity loss or broadcasting separability/entanglement loss.

3. Loss Functions and Evaluation Metrics

a. Loss Functions: To ensure the generator produces physically valid and task-specific quantum states, we employ a composite loss function, G_{loss} . This function combines the standard adversarial loss with several penalty terms designed to enforce the required quantum properties. Key components include penalties for deviations from positive semi-definiteness (L_{psd}) and the trace-one condition (L_{trace}). Furthermore, we introduce a utility-based loss, L_{task} , which for the teleportation task specifically penalizes states that do not meet the required fidelity threshold for a quantum advantage. The full mathematical formulation of each loss component and the final weighted-sum equation are detailed in Appendix B.

b. Evaluation Metrics: To assess performance, we use three primary metrics to compare the set of generated states, \mathcal{G} , against the set of training states, \mathcal{T} .

b1. Accuracy: We define accuracy as the fraction of generated states that satisfy the predefined utility criteria for the target resource class (e.g., $F_{\text{max}} > 2/3$). This metric directly evaluates the generator's success in producing useful states.

b2. Average Cross-Set Fidelity: To measure state-level similarity, we compute the average fidelity between the set of generated states \mathcal{G} and the set of training states \mathcal{T} :

$$F(\mathcal{G}, \mathcal{T}) = \frac{1}{|\mathcal{G}||\mathcal{T}|} \sum_{\rho_g \in \mathcal{G}} \sum_{\rho_t \in \mathcal{T}} F(\rho_g, \rho_t), \quad (10)$$

where $F(\rho_g, \rho_t)$ is the Uhlmann fidelity between two density matrices. A score closer to the average pairwise

fidelity of the training data with itself indicates a perfect match in both quality and variety. This intra-dataset fidelity serves as an ideal benchmark, and we plot it as a horizontal baseline in our results to provide a clear target for the generator's performance.

b3. Fréchet Inception Distance (FID): To assess distributional similarity between generated and training states, we adapt the Fréchet Inception Distance [70]. Feature embeddings $\phi(\rho)$ are computed as Pauli-expectation vectors: for a two-qubit state, $\phi(\rho) = (\langle \sigma_i \otimes \sigma_j \rangle)_{i,j=0}^3$, yielding a 16-dimensional real vector. The training and generated feature distributions are modeled as multivariate Gaussians, with means $(\mu_{\mathcal{G}}, \mu_{\mathcal{T}})$ and covariances $(\Sigma_{\mathcal{G}}, \Sigma_{\mathcal{T}})$. The FID score is then given by Eq. (11):

$$\begin{aligned} \text{FID}(\mathcal{G}, \mathcal{T}) = & \|\mu_{\mathcal{G}} - \mu_{\mathcal{T}}\|^2 \\ & + \text{Tr}\left(\Sigma_{\mathcal{G}} + \Sigma_{\mathcal{T}} - 2\sqrt{\Sigma_{\mathcal{G}}\Sigma_{\mathcal{T}}}\right). \end{aligned} \quad (11)$$

Lower values indicate that the generator not only produces valid states but also captures the diversity and covariance structure of the target distribution. Unlike average fidelity, which measures pairwise similarity at the state level, FID is sensitive to coverage and mode collapse, providing complementary insight into how well the GAN reproduces the overall training distribution.

IV. GENERATION OF RESOURCE STATES THROUGH CGAN

Having established the CGAN framework and evaluation methodology, we now apply it to generate quantum resource states for broadcasting and teleportation. In this section, we discuss the generation of two different types of resource states (broadcastable and teleportable) through the three generator architectures described above. We do these for two different types of states, namely Bell diagonal and Werner like states. In the later part of the section we also compare the states generated with the known theoretical ranges for these classes of states. Before we go to the first subsection where we analyze and compare between the generation process, we give a brief description of the inputs that we have taken.

a. Werner-like states. These states are defined by the Bloch vector and correlation matrix given in Eq. (12):

$$\begin{aligned} \rho_{12}^w = & \{\vec{x}^w, \vec{x}^w, T^w\}, \\ \vec{x}^w = & \{0, 0, p(\alpha^2 - \beta^2)\}, \\ T^w = & \text{diag}(2p\alpha\beta, -2p\alpha\beta, p), \end{aligned} \quad (12)$$

with the conditions $\alpha^2 + \beta^2 = 1$ and $0 \leq p \leq 1$. Note that the diagonal matrix $M = \text{diag}(\dots)$ has diagonal elements specified inside the brackets.

b. Bell-diagonal states. These states can be expressed via Eq. (13), where the Bloch vector is null:

$$\begin{aligned} \rho_{12}^b = & \{\vec{0}, \vec{0}, T^b\}, \\ T^b = & \text{diag}(c_1, c_2, c_3), \end{aligned} \quad (13)$$

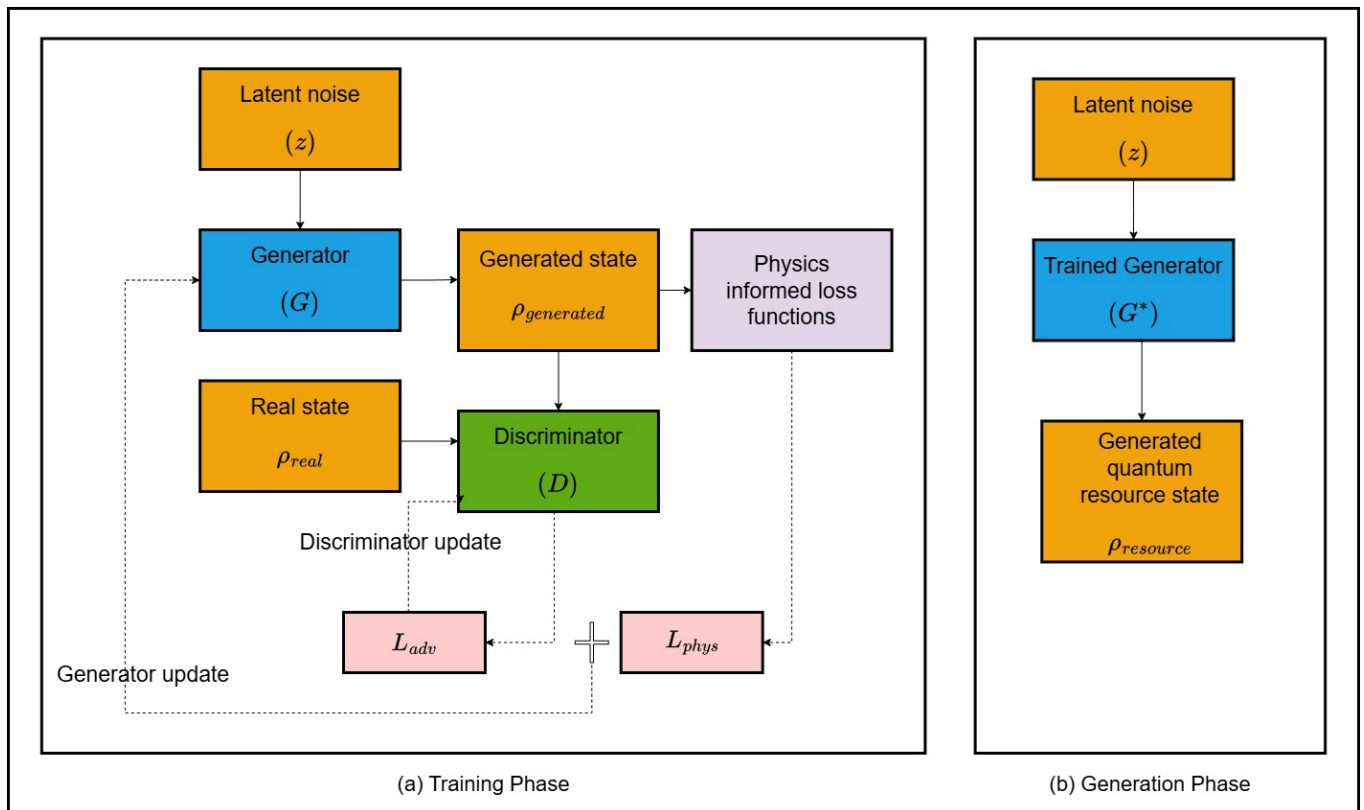


FIG. 1. Physics-informed CGAN framework used in this work. (a) Training: the generator G maps latent noise z to ρ_{gen} . The discriminator D compares ρ_{gen} with training samples ρ_{train} , producing the adversarial term L_{adv} . In parallel, a physics module computes constraint and utility terms (trace-one, PSD, and task utility), yielding the weighted objective in Eq. (9). The discriminator and generator are updated alternately. (b) Inference: after training, samples z are mapped by the trained generator G^* to ρ_{resource} , producing candidate quantum resource states.

with $-1 \leq c_i \leq 1$. The eigenvalues of ρ_{12}^b are given by Eq. (14):

$$\lambda_{mn} = \frac{1}{4} [1 + (-1)^m c_1 - (-1)^{m+n} c_2 + (-1)^n c_3], \quad (14)$$

A. Generation and Analysis

In the following analysis, Fig. 2 reports task accuracy, Fig. 3 reports average cross-set fidelity, and Fig. 4 reports Fréchet Inception Distance (FID). Each figure presents (1) Bell-diagonal and (2) Werner-like states, with three subplots per row: (a) local broadcasting, (b) non-local broadcasting, and (c) teleportation. Complete numerical results are tabulated in Appendix C.

Across all experiments, a clear distinction emerged between constraint-enforcing and unconstrained generator architectures. Decomposition-based models (Cholesky and Enhanced LDL) consistently achieved rapid convergence and stable accuracy, with average cross-set fidelity scores closely tracking theoretical benchmarks for both Werner-like and Bell-diagonal states. These models, by

construction, maintained physical validity throughout training, reducing the need for heavy penalties for auxiliary losses. However, their sensitivity to hyperparameter tuning also introduced risks of mode collapse in smaller training sets.

In contrast, the direct generation model demonstrated greater robustness to hyperparameter variation but consistently lagged behind decomposition-based approaches in distribution-level measures such as Fréchet Inception Distance (FID) across all studied tasks. This indicates that while unconstrained models offer flexibility and adaptability, they did not surpass decomposition models in FID within the two-qubit scenarios considered here.

This performance gap is a key finding. It demonstrates that for the well-structured two-qubit problems studied here, embedding physical constraints directly into the generator's architecture is a far more effective and stable strategy than relying on penalty-based learning. While the flexibility of the direct model remains appealing for future, less-understood problems, our results underscore the significant value of architecturally-enforced constraints when the underlying physics is known.

Comparison with theoretical boundaries further validates

the generative process. Scatter plots of generated Bell-diagonal states (Fig. 5) confirmed confinement to the appropriate regions, with the brown wireframe depicting the octahedron of inseparable states and the teal shaded areas marking the resource subregions at each vertex for broadcastable and teleportable states. Similarly, the Werner-like states generated (Fig. 6) adhered to the theoretical boundaries, with the overwhelming majority of useful states (blue dots) falling within the green entangled region. The small number of orange dots appearing in the green region represent minor classification inaccuracies, while the overall distribution strongly reproduces the known analytical boundary $p(1 + 4\alpha\beta) > 1$ for teleportability (see Appendix D 2). These findings establish that the proposed CGAN framework does not merely memorize training samples but generalizes to produce physically valid states consistent with analytic criteria.

Figure 2 summarizes the task-accuracy trajectories during training. Across both state families, the constraint-enforcing generators (Cholesky Decomposition and Enhanced LDL Decomposition) reach high accuracy early and remain comparatively stable across epochs, indicating that once the generator learns to stay within the physically valid manifold, satisfying the task utility condition becomes reliably achievable. In contrast, the Direct Matrix Generator exhibits substantially lower accuracy across tasks and a noticeably slower rise, consistent with the additional burden of learning physical validity and task structure purely through penalty terms.

A second qualitative trend visible in Fig. 2 is task dependence: broadcasting tasks exhibit greater variability across epochs than teleportation, which is expected because the broadcasting labels are defined via separability/entanglement constraints on cloned output marginals rather than a single scalar threshold. For both Bell-diagonal and Werner-like settings, the decomposition-based models maintain accuracy curves that are consistently closer to the desired accuracy than the direct model, supporting the conclusion that architectural constraint enforcement improves stability for these two-qubit resource-generation tasks.

Figure 3 compares the models on average cross-set fidelity, which measures how well the distribution of generated states matches the training data. For both broadcastable and teleportable Werner-like states, the decomposition-based *Cholesky Decomposition* and *Enhanced LDL Decomposition* models perform exceptionally well, achieving and maintaining high fidelity. The *Direct Matrix Generator*, in contrast, shows a significant drop in fidelity, indicating a failure to correctly learn the target distribution for these more complex states.

In Fig. 3, the dashed horizontal baseline denotes the average fidelity of the training set with itself, which serves as an internal reference for the typical pairwise similarity scale within the target distribution. Fidelity curves that approach this baseline indicate that the generator is pro-

ducing states that are statistically similar to the training data at the state level, while sustained gaps indicate a mismatch in either typical state quality or coverage.

Qualitatively, the decomposition-based models track the training baseline more closely across both Bell-diagonal and Werner-like tasks, with particularly strong agreement in the Werner-like broadcasting cases where the curves remain high and stable. By comparison, the Direct Matrix Generator shows larger deviations in several tasks, consistent with reduced agreement at the state level. These observations align with the accuracy trends: models that enforce positivity and trace structurally not only satisfy the utility criteria more often, but also tend to generate states that remain closer to the training distribution under the fidelity metric.

Finally, Figure 4 presents Fréchet Inception Distance (FID) between generated and training distributions. Decomposition-based models consistently achieve lower final FID scores, indicating closer alignment with training sets. The Direct Matrix Generator sometimes exhibits visually stable trajectories (notably in nonlocal Bell-diagonal and teleportable Werner-like tasks), but its reported endpoint FID values remain higher, confirming the quantitative advantage of decomposition models. The Direct Matrix Generator also exhibits FID instability in specific cases: for Werner-like local broadcasting, FID diverges mid-training and fails to recover, while for Bell-diagonal local broadcasting, FID increases sharply towards the end of training. This instability likely stems from violations of positive semidefiniteness that distort the feature covariance.

FID provides a complementary, distribution-level view: lower values indicate that the generated set reproduces not only typical samples but also the mean and covariance structure of the feature representation used (Pauli-expectation embeddings). Thus, FID is particularly sensitive to reduced diversity or partial mode collapse, even in cases where pointwise fidelity can remain moderate.

The trajectories in Fig. 4 show that decomposition-based generators typically converge to lower final FID values than the Direct Matrix Generator across tasks, indicating closer alignment with the target distribution at the level of first- and second-order statistics. In several settings the direct model exhibits comparatively flatter or slowly improving curves, but with higher endpoints, consistent with a residual distribution mismatch. Taken together with Fig. 3, these trends support the interpretation that enforcing physical structure in the generator improves both task success and distributional coverage for the two-qubit cases studied here.

B. Comparison with Theoretical Results

Figure 5 shows Bell-diagonal states generated and plotted using the trained models. For Bell-diagonal states, when




FIG. 2. Accuracy comparison across generator architectures (training size = 500). (1) Bell-diagonal states; (2) Werner-like states. Subplots: (a) local broadcasting, (b) non-local broadcasting, (c) teleportation.

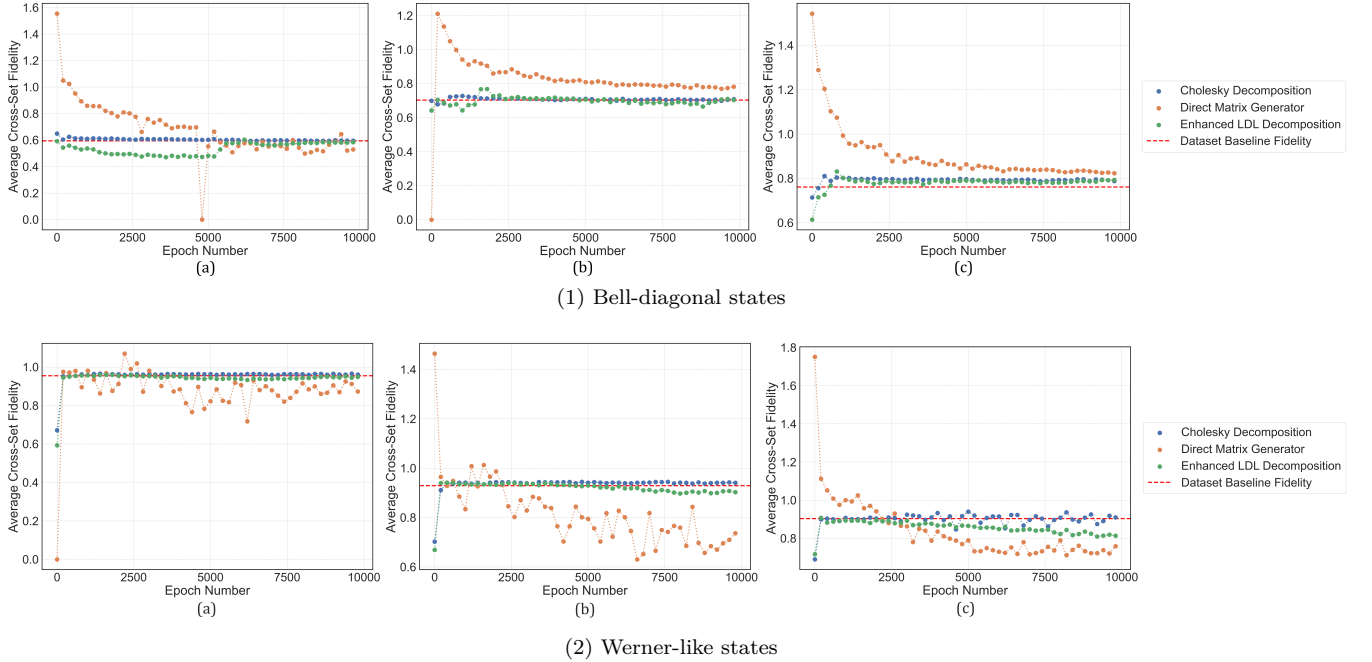


FIG. 3. Average cross-set fidelity between generated and training states (training size = 1000). (1) Bell-diagonal states; (2) Werner-like states. Subplots: (a) local broadcasting, (b) non-local broadcasting, (c) teleportation. The dashed red line indicates the ideal benchmark—the average fidelity of the training data with itself.

we plot c_1, c_2, c_3 , all valid states lie within a tetrahedron with vertices: A : $(1, -1, 1)$, B : $(-1, 1, 1)$, C : $(1, 1, -1)$, D : $(-1, -1, -1)$. In these plots, the brown wireframe depicts the boundary of the octahedron of inseparable (entangled) Bell-diagonal states, which lies inside the larger tetrahedron of all valid Bell-diagonal states. The teal

shaded subregions near each vertex indicate the theoretical parameter ranges where states serve as resources for the respective task. Generated states classified as useful are shown in green, while those not meeting the resource criterion are shown in red.

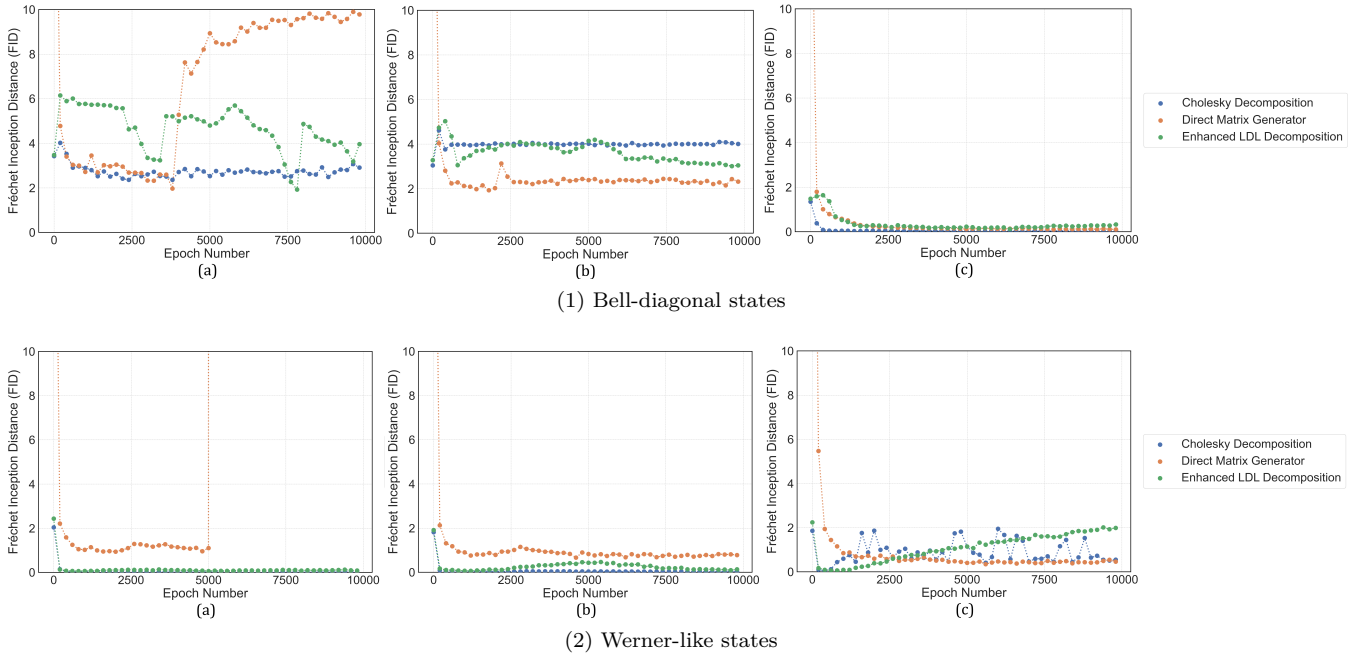


FIG. 4. Fréchet Inception Distance (FID) between generated and training state distributions (training size = 2000). (1) Bell-diagonal states; (2) Werner-like states. Subplots: (a) local broadcasting, (b) non-local broadcasting, (c) teleportation. Lower FID indicates closer distributional match.

All three subplots of Figure 5 show states generated using Model 2 (Enhanced LDL) with training size 2000.

(1) *Local broadcasting*: States useful for broadcasting of quantum correlations using local cloning, concentrated around the tetrahedron vertices. The teal region corresponds to the following vertex coordinates:

$$\begin{aligned}
 D: & (-1, -1, -1), \left(-\frac{5}{8}, -\frac{5}{8}, -1\right), \left(-1, -\frac{5}{8}, -\frac{5}{8}\right), \left(-\frac{5}{8}, -1, -\frac{5}{8}\right) \\
 C: & (1, 1, -1), \left(\frac{5}{8}, \frac{5}{8}, -1\right), \left(1, \frac{5}{8}, -\frac{5}{8}\right), \left(\frac{5}{8}, 1, -\frac{5}{8}\right) \\
 B: & (-1, 1, 1), \left(-\frac{5}{8}, \frac{5}{8}, 1\right), \left(-1, \frac{5}{8}, \frac{5}{8}\right), \left(-\frac{5}{8}, 1, \frac{5}{8}\right) \\
 A: & (1, -1, 1), \left(\frac{5}{8}, -\frac{5}{8}, 1\right), \left(1, -\frac{5}{8}, \frac{5}{8}\right), \left(\frac{5}{8}, -1, \frac{5}{8}\right)
 \end{aligned}$$

(2) *Non-local broadcasting*: States useful for broadcasting of quantum correlations using non-local cloning, spanning a larger area around the vertices compared to local broadcasting. The teal region corresponds to:

$$\begin{aligned}
 D: & (-1, -1, -1), \left(-\frac{1}{3}, -\frac{1}{3}, -1\right), \left(-1, -\frac{1}{3}, -\frac{1}{3}\right), \left(-\frac{1}{3}, -1, -\frac{1}{3}\right) \\
 C: & (1, 1, -1), \left(\frac{1}{3}, \frac{1}{3}, -1\right), \left(1, \frac{1}{3}, -\frac{1}{3}\right), \left(\frac{1}{3}, 1, -\frac{1}{3}\right) \\
 B: & (-1, 1, 1), \left(-\frac{1}{3}, \frac{1}{3}, 1\right), \left(-1, \frac{1}{3}, \frac{1}{3}\right), \left(-\frac{1}{3}, 1, \frac{1}{3}\right) \\
 A: & (1, -1, 1), \left(\frac{1}{3}, -\frac{1}{3}, 1\right), \left(1, -\frac{1}{3}, \frac{1}{3}\right), \left(\frac{1}{3}, -1, \frac{1}{3}\right)
 \end{aligned}$$

(3) *Teleportation*: States useful for teleportation, spanning specific regions within the tetrahedron:

$$\begin{aligned}
 D: & (-1, -1, -1), (-1, 0, 0), (0, -1, 0), (0, 0, -1) \\
 C: & (1, 1, -1), (0, 1, 0), (1, 0, 0), (0, 0, -1) \\
 B: & (-1, 1, 1), (-1, 0, 0), (0, 1, 0), (0, 0, 1) \\
 A: & (1, -1, 1), (0, -1, 0), (0, 0, 1), (1, 0, 0)
 \end{aligned}$$

Figure 6 presents Werner-like states generated and plotted using the trained models. In these plots, the pink background indicates the region of separable states, while the green background marks the region of inseparable (entangled) states. For the broadcasting tasks (1) and (2), the green region represents all entangled states from which broadcasting is possible; for teleportation (3), the green region coincides with states useful for teleportation. Generated states that meet the resource criterion are shown as blue dots, while those that do not are shown as orange dots; states lying in the “wrong” background region thus represent classification errors by the model.

All three subplots show states generated using Model 1 (Cholesky) with training size 2000: (1) states useful for broadcasting of quantum correlations using local cloning, (2) states useful for broadcasting of quantum correlations using non-local cloning, and (3) states useful for teleportation, satisfying the utility criterion $p(1 + 4\alpha\beta) > 1$.

V. DISCUSSION AND CONCLUSION

The results presented establish a systematic method for generating quantum states tailored to specific tasks. In this work we have chosen these tasks to be broadcasting of quantum correlation and teleportation. The central finding is the clear performance gap between architecturally constrained generators (Cholesky, LDL) and an unconstrained, direct-output model. For the well-defined two-




FIG. 5. Generated Bell-diagonal states plotted in the parameter space (c_1, c_2, c_3) . Subplots: (1) local broadcasting, (2) non-local broadcasting, (3) teleportation. The brown wireframe outlines the octahedron of inseparable (entangled) Bell-diagonal states, contained within the larger tetrahedron of all valid Bell-diagonal states. The teal shaded subregions indicate the theoretical parameter ranges where states are useful for each task. Green dots represent generated states meeting the resource criterion; red dots represent states not meeting it. All generated states are from models trained with dataset size 2000.




FIG. 6. Generated Werner-like states plotted in terms of the parameters p and α . Subplots: (1) local broadcasting, (2) non-local broadcasting, (3) teleportation. The green background denotes the region of inseparable (entangled) states; the pink background indicates the separable region. For (1) and (2), the green region represents all entangled states from which broadcasting is possible; for (3), it coincides with states useful for teleportation. Blue dots represent generated states classified as useful; orange dots represent those not meeting the criterion. All generated states are from models trained with dataset size 2000.

qubit systems studied, embedding physical constraints directly into the generator’s design proved to be a more effective and stable learning strategy than relying on penalty-based loss functions to enforce validity. This suggests that for problems where the underlying physical structure is known, leveraging that structure within the model architecture is paramount.

Furthermore, our framework’s ability to reproduce the known theoretical boundaries for teleportable and broadcastable states (Figs. 5 and 6) serves as a critical validation. The visual alignment of generated states with theoretical subregions—bounded by the brown wireframe

octahedron of inseparable states for Bell-diagonal states and the entanglement boundary for Werner-like states—confirms that the CGAN does not merely memorize the training data but learns a generalized representation of the target resource state manifold. The direct generator, while less accurate, demonstrated robustness to hyperparameter variations, indicating its potential utility in more exploratory scenarios where physical constraints are less understood or analytically intractable, and where flexibility is more valuable than guaranteed precision. Our results is about generating task specific quantum resources in particular to the context of quantum communication

which can act as the backbone of generation of resourceful network [34, 50].

The present study is restricted to two-qubit systems and two specific state families (Werner-like and Bell-diagonal), for which analytical criteria exist. Whether the performance advantage of decomposition-based generators persists for higher-dimensional or less-structured state spaces remains an open question. Additionally, the framework assumes access to a well-defined utility function; tasks lacking computable metrics would require alternative formulations. These observations motivate several directions for extending and refining the framework.

VI. FUTURE OUTLOOK

Our results establish that physics-informed adversarial learning can generate families of valid quantum resource states tuned to protocol-specific utilities. The immediate challenge is to extend this approach beyond two-qubit systems. While computational scaling is favorable ($O(d^3)$ for validation), defining computable utilities for multipartite

entanglement [31], quantum metrology [71], and error correction [72] remains non-trivial. A natural extension is to replace the task-specific loss with other computable metrics—for instance, Quantum Fisher Information for metrology or the Geometric Measure of Entanglement for multipartite state discovery—enabling generative exploration where closed-form criteria are unknown.

A complementary direction is architectural: hybrid generators that couple decomposition-based structure with learnable residual components may balance physical faithfulness and expressivity. Embedding known symmetries and conservation laws [73] could further reduce sample complexity. Finally, interfacing the framework with quantum-network simulators such as NetSquid [74] would permit testing under realistic noise and topology models, providing a bridge between numerical design and near-term experiments.

Acknowledgment: I.C. acknowledges the support from the Ministry of Electronics and Information Technology (MeitY), Government of India, under Grant No. 4(3)/2024-ITEA. I.C. thanks Soumit Roy for useful discussions.

Data Availability : [GitHub Repository for Project](#)

-
- [1] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade, *Photonic Crystals: Molding the Flow of Light*, 2nd ed. (Princeton University Press, 2008).
 - [2] D. R. Smith, J. B. Pendry, and M. C. Wiltshire, *Science* **305**, 788 (2004).
 - [3] S. Molesky, Z. Lin, A. Y. Piggott, W. Jin, J. Vuckovic, and A. W. Rodriguez, *Nat. Photon.* **12**, 659 (2018).
 - [4] I. Malkiel, M. Mrejen, A. Nagler, U. Arieli, L. Wolf, and H. Suchowski, *Light Sci. Appl.* **7**, 60 (2018).
 - [5] S. Zhao, G. Lee, J. Li, and X. Hu, *APL Mater.* **7**, 101103 (2019).
 - [6] J. Schmidt, M. R. Marques, S. Botti, and M. A. Marques, *npj Comput. Mater.* **5**, 83 (2019).
 - [7] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, *Rev. Mod. Phys.* **91**, 045002 (2019).
 - [8] J. Shlomi, P. Battaglia, and J.-R. Vlimant, *Mach. Learn.: Sci. Technol.* **2**, 021001 (2020).
 - [9] M. A. Aragon-Calvo, *Mon. Not. R. Astron. Soc.* **484**, 5771 (2019).
 - [10] J. Carrasquilla, *Adv. Phys.: X* **5**, 1797528 (2020).
 - [11] R. Iten, T. Metger, H. Wilming, L. Del Rio, and R. Renner, *Phys. Rev. Lett.* **124**, 010508 (2020).
 - [12] G. P. Zhang, *IEEE Trans. Syst. Man Cybern. Part C Appl. Rev.* **30**, 451 (2002).
 - [13] D. P. Kingma and M. Welling, arXiv preprint arXiv:1312.6114 (2013).
 - [14] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, *Advances in neural information processing systems* **27** (2014).
 - [15] S. Ahmed, C. Sánchez Muñoz, F. Nori, and A. F. Kockum, *Phys. Rev. Lett.* **127**, 140502 (2021).
 - [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, *Commun. ACM* **63**, 139–144 (2020).
 - [17] K. Bharti, T. Haug, V. Vedral, and L.-C. Kwek, *AVS Quant. Sci.* **2**, 10.1116/5.0007529 (2020).
 - [18] B. S. Rem, N. Käming, M. Tarnowski, L. Asteria, N. Fläschner, C. Becker, K. Sengstock, and C. Weitenberg, *Nat. Phys.* **15**, 917 (2019).
 - [19] Y.-C. Ma and M.-H. Yung, *npj Quant. Inf.* **4**, 34 (2018).
 - [20] C. Harney, S. Pirandola, A. Ferraro, and M. Paternostro, *N. J. Phys.* **22**, 045001 (2020).
 - [21] G. Sentís, M. Guță, and G. Adesso, *EPJ Quant. Tech.* **2**, 1 (2015).
 - [22] V. Gebhart and M. Bohmann, *Phys. Rev. Res.* **2**, 023150 (2020).
 - [23] M. Krenn, M. Erhard, and A. Zeilinger, *Nat. Rev. Phys.* **2**, 649 (2020).
 - [24] G. Torlai and R. G. Melko, *Phys. Rev. Lett.* **119**, 030501 (2017).
 - [25] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, *Phys. Rev. X* **8**, 031084 (2018).
 - [26] N. Wittler, F. Roy, K. Pack, M. Werninghaus, A. S. Roy, D. J. Egger, S. Filipp, F. K. Wilhelm, and S. Machnes, *Phys. Rev. App.* **15**, 034080 (2021).
 - [27] E. Chitambar and G. Gour, *Rev. Mod. Phys.* **91**, 025001 (2019).
 - [28] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge university press, 2010).
 - [29] R. Horodecki, M. Horodecki, and P. Horodecki, *Phys. Lett. A* **222**, 21 (1996).
 - [30] S. Chatterjee, S. Sazim, and I. Chakrabarty, *Phys. Rev. A* **93**, 042309 (2016).
 - [31] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, *Rev. Mod. Phys.* **81**, 865 (2009).

- [32] N. Ganguly, S. Adhikari, A. S. Majumdar, and J. Chatterjee, *Phys. Rev. Lett.* **107**, 270501 (2011).
- [33] M. Vempati, N. Ganguly, I. Chakrabarty, and A. K. Pati, *Phys. Rev. A* **104**, 012417 (2021).
- [34] G. Mylavarapu, S. Ghosh, C. Hens, I. Chakrabarty, and S. Mitra, arXiv:2409.20304 [10.48550/arXiv.2409.20304](https://arxiv.org/abs/2409.20304) (2024).
- [35] S. Patro, I. Chakrabarty, and N. Ganguly, *Phys. Rev. A* **96**, 062102 (2017).
- [36] P. V. Srinidhi, I. Chakrabarty, S. Bhattacharya, and N. Ganguly, *Phys. Rev. A* **110**, 042423 (2024).
- [37] M. Vempati, S. Shah, N. Ganguly, and I. Chakrabarty, *Quantum* **6**, 641 (2022).
- [38] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, *Phys. Rev. Lett.* **70**, 1895 (1993).
- [39] S. Wehner, D. Elkouss, and R. Hanson, *Science* **362**, eaam9288 (2018).
- [40] C. H. Bennett and S. J. Wiesner, *Phys. Rev. Lett.* **69**, 2881 (1992).
- [41] P. Horodecki and R. Horodecki, *Quantum Inf. Comput.* **1**, 45 (2001).
- [42] P. Singh and I. Chakrabarty, *Phys. Rev. A* **109**, 032406 (2024).
- [43] W. Hu, R. R. Singh, and R. T. Scalettar, *Phys. Rev. E* **95**, 062122 (2017).
- [44] M. Paganini, L. de Oliveira, and B. Nachman, *Phys. Rev. Lett.* **120**, 042003 (2018).
- [45] S. Lloyd and C. Weedbrook, *Phys. Rev. Lett.* **121**, 040502 (2018).
- [46] L. Hu, S.-H. Wu, W. Cai, Y. Ma, X. Mu, Y. Xu, H. Wang, Y. Song, D.-L. Deng, C.-L. Zou, and L. Sun, *Sci. Adv.* **5**, eaav2761 (2019).
- [47] D. Koutný, L. Motka, Z. Hradil, J. Řeháček, and L. L. Sánchez-Soto, *Phys. Rev. A* **106**, 012409 (2022).
- [48] J. Ho, A. Jain, and P. Abbeel, in *Advances in Neural Information Processing Systems*, Vol. 33 (2020) pp. 6840–6851.
- [49] B. Sanchez-Lengeling and A. Aspuru-Guzik, *Science* **361**, 360 (2018).
- [50] S. Roy, M. R. Miraj, C. Hens, G. Mylavarapu, S. Ghosh, and I. Chakrabarty, arXiv preprint arXiv:2508.10417 (2025).
- [51] V. Buzek, V. Vedral, M. B. Plenio, P. L. Knight, and M. Hillery, *Phys. Rev. A* **55**, 3327 (1997).
- [52] R. Mundra, D. Patel, I. Chakrabarty, N. Ganguly, and S. Chatterjee, *Phys. Rev. A* **100**, 042319 (2019).
- [53] U. K. Sharma, I. Chakrabarty, and M. K. Shukla, *Phys. Rev. A* **96**, 052319 (2017).
- [54] D. Patel, A. Roy, I. Chakrabarty, and N. Ganguly, *Pramana* **98**, 1 (2024).
- [55] K. Modi, T. Paterek, W. Son, V. Vedral, and M. Williamson, *Phys. Rev. Lett.* **104**, 080501 (2010).
- [56] A. Jain, I. Chakrabarty, and S. Chatterjee, *Phys. Rev. A* **99**, 022315 (2019).
- [57] A. Peres, *Quantum theory: concepts and methods*, Vol. 72 (Springer, 1997).
- [58] R. F. Werner, *Phys. Rev. A* **40**, 4277 (1989).
- [59] R. Mundra, S. Chattopadhyay, I. Chakrabarty, and N. Ganguly, *Pramana* **95**, 60 (2021).
- [60] S. Adhikari, A. S. Majumdar, and N. Nayak, *Phys. Rev. A* **77**, 042301 (2008).
- [61] H. Barnum, J. Barrett, M. Leifer, and A. Wilce, *Phys. Rev. Lett.* **99**, 240501 (2007).
- [62] W. K. Wootters and W. H. Zurek, *Nature* **299**, 802 (1982).
- [63] I. Chakrabarty, *Eur. Phys. J. D* **57**, 265 (2010).
- [64] I. Chakrabarty, N. Ganguly, and B. S. Choudhury, *Quantum Inf. Process.* **10**, 27 (2011).
- [65] S. Adhikari, N. Ganguly, I. Chakrabarty, and B. S. Choudhury, *J. Phys. A: Math. Theor.* **41**, 415302 (2008).
- [66] M. Horodecki, P. Horodecki, and R. Horodecki, *Phys. Rev. Lett.* **80**, 5239 (1998).
- [67] G. Mylavarapu, I. Chakrabarty, K. Mukherjee, M. Huang, and J. Wu, arXiv:2306.01406 [10.48550/arXiv.2306.01406](https://arxiv.org/abs/2306.01406) (2023).
- [68] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, in *Advances in Neural Information Processing Systems*, Vol. 27, edited by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Curran Associates, Inc., 2014) pp. 2672–2680.
- [69] M. Mirza and S. Osindero, arXiv preprint arXiv:1411.1784 [10.48550/arXiv.1411.1784](https://arxiv.org/abs/1411.1784) (2014).
- [70] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, in *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17* (Curran Associates Inc., Red Hook, NY, USA, 2017) p. 6629–6640.
- [71] V. Giovannetti, S. Lloyd, and L. Maccone, *Nat. Photon.* **5**, 222 (2011).
- [72] B. M. Terhal, *Rev. Mod. Phys.* **87**, 307 (2015).
- [73] T. S. Cohen, M. Weiler, B. Kicanaoglu, and M. Welling, in *Proceedings of the 36th International Conference on Machine Learning*, PMLR, Vol. 97 (2019) pp. 1321–1330.
- [74] T. Coopmans, R. Knegjens, A. Dahlberg, D. Maier, L. Nijsten, J. de Oliveira Filho, M. Papendrecht, J. Rabbie, F. Rozpedek, M. Skrzypczyk, L. Wubben, W. de Jong, D. Podberšek, A. Torres-Knoop, D. Elkouss, and S. Wehner, *Commun. Phys.* **4**, 164 (2021).

Appendix A: Reproducibility Details

This section contains all necessary information for the full reproduction of the results presented in the main text.

1. Model Architectures

All generator models take a 100-dimensional latent vector $z \sim \mathcal{N}(0, I)$ and include a residual connection for stability. A single discriminator architecture is shared across experiments.

a. Generator models. Cholesky Decomposition: Parameterizes $\rho = LL^\dagger$, where any matrix of the form LL^\dagger is inherently positive semi-definite, enforcing this constraint by construction. Uses a multilayer perceptron (MLP) with Leaky ReLU activations, LayerNorm, and a trace normalization step.

Enhanced LDL Decomposition: Uses $\rho = LDL^\dagger$ with L lower triangular (unit diagonal) and D diagonal with strictly positive entries ($\epsilon = 10^{-6}$ regularization). Ensures Hermiticity by construction and trace-one via normalization.

Direct Matrix Generator: Outputs a 32-dimensional vector reshaped into a 4×4 Hermitian matrix. Physical constraints (Hermiticity, PSD, trace-one) are enforced entirely through explicit loss functions rather than architectural design, serving as a baseline to demonstrate the benefit of constraint embedding.

b. Discriminator. An MLP classifier for flattened 4×4 matrices (32 real values): 1. Input: Linear(32 \rightarrow 256) + LeakyReLU + LayerNorm + Dropout. 2. Hidden Layers: 256 \rightarrow 512, 512 \rightarrow 256, 256 \rightarrow 128.3. Output: Linear(128 \rightarrow 1) with Sigmoid

2. Training Hyperparameters

All models were trained with the RMSprop optimizer. Loss weights were tuned empirically to balance contributions.

TABLE A1. Training hyperparameters and loss weights for decomposition-based generators (Models 1 & 2).

Hyperparameter	Value
Optimizer	RMSprop
Learning Rate	10^{-5}
Batch Size	512
Training Epochs	10,000
λ_{psd} (PSD Loss)	10.0
λ_{trace} (Trace Loss)	10.0
λ_{herm} (Hermiticity Loss, Model 3 only)	5.0
λ_{task} (Task-Specific Loss)	$5.0 \times m_{task}$
$\lambda_{diversity}$ (Diversity Loss)	0.5

The task-specific multiplier m_{task} accounts for varying optimization complexity: Teleportation: $m_{task} = 1.0$ (simpler $F_{max} > 2/3$ criterion). Nonlocal Broadcasting: $m_{task} = 1.2$. Local Broadcasting: $m_{task} = 1.5$ (most complex Positive Partial Transpose (PPT) conditions).

a. Rationale. For decomposition-based models (Cholesky/LDL), physics constraints (PSD, trace-one) are architecturally enforced, reducing the need for strong penalty terms. The moderate physics weights provide soft regularization while the elevated task-specific weight ensures the optimization focuses on the actual objective. The diversity term encourages exploration without dominating convergence. RMSprop was selected as the default optimizer after comparing Adadelta, AdamW, Adam, and RMSprop; it showed greater stability on constrained broadcasting tasks.

3. Dataset Generation

The training datasets used for the discriminator were generated as follows:

1. *Werner-like states:* Training data were generated by uniformly sampling the parameters $p \in [0, 1]$ and α, β

from a uniform distribution on the unit circle (subject to $\alpha^2 + \beta^2 = 1$), retaining only states satisfying the task-specific criterion (e.g., $p(1 + 4\alpha\beta) > 1$ for teleportation; see Appendix D 2). For each task, 2000 unique states were generated.

2. *Bell-diagonal states:* Training data were generated by uniformly sampling the parameters c_1, c_2, c_3 from the range $[-1, 1]$, subject to the constraints required for the state to be a valid resource for the specified task (e.g., $|c_1| + |c_2| + |c_3| > 1$ for teleportation). For each task, 2000 unique states were generated.

Appendix B: Loss Function Formulation

This section provides the explicit mathematical forms of the loss components referenced in Eq. (9) of the main text. For a batch of N generated density matrices $\{\rho_{gen}^{(i)}\}_{i=1}^N$, the individual components are:

- *Adversarial loss:*

$$\mathcal{L}_{adv} = -\frac{1}{N} \sum_{i=1}^N \log(D(\rho_{gen}^{(i)}))$$

- *Trace loss:*

$$\mathcal{L}_{trace} = \frac{1}{N} \sum_{i=1}^N \left| \text{Tr}(\rho_{gen}^{(i)}) - 1 \right|$$

- *PSD loss:* For eigenvalues $\{\lambda_j^{(i)}\}$ of each $\rho_{gen}^{(i)}$:

$$\mathcal{L}_{psd} = \frac{1}{N} \sum_{i=1}^N \sum_j \max(0, -\lambda_j^{(i)})$$

- *Task-specific loss (\mathcal{L}_{task}):* Penalizes states failing the fidelity criterion— $F_{max} \leq 2/3$ for teleportation, or PPT-based conditions for broadcasting.

Appendix C: Experimental Results

Tables C1–C6 present detailed performance metrics across models, tasks, and dataset sizes. Abbreviated model names: Cholesky (Cholesky Decomposition), LDL (Enhanced LDL Decomposition), Direct (Direct Matrix Generator). Analysis of these results is provided in Section IV of the main text.

Note: Metric definitions are provided in Section III of the main text.

Computational Complexity of State Generation

To understand the scaling behavior of our generator models and to contextualize the empirical benchmarks pre-

TABLE C1. Metrics for Bell Diagonal — Local Broadcasting.

Model	Train Size	Accuracy	Fidelity	FID
Cholesky	500	0.936	0.590	0.517
Cholesky	1000	0.880	0.586	0.451
Cholesky	2000	0.941	0.581	0.452
LDL	500	0.988	0.483	0.885
LDL	1000	0.679	0.593	0.467
LDL	2000	0.652	0.659	0.560
Direct	500	0.100	0.681	0.678
Direct	1000	0.120	0.514	2.236
Direct	2000	0.150	0.467	2.394

TABLE C2. Metrics for Bell Diagonal — Nonlocal Broadcasting.

Model	Train Size	Accuracy	Fidelity	FID
Cholesky	500	0.881	0.710	0.083
Cholesky	1000	0.876	0.700	0.049
Cholesky	2000	0.950	0.701	0.123
LDL	500	0.983	0.678	0.121
LDL	1000	0.770	0.694	0.127
LDL	2000	0.982	0.704	0.148
Direct	500	0.120	0.744	0.384
Direct	1000	0.150	0.703	0.489
Direct	2000	0.180	0.702	0.442

TABLE C3. Metrics for Bell Diagonal — Teleportable.

Model	Train Size	Accuracy	Fidelity	FID
Cholesky	500	0.743	0.790	0.011
Cholesky	1000	0.852	0.789	0.016
Cholesky	2000	0.979	0.769	0.011
LDL	500	0.980	0.798	0.010
LDL	1000	0.927	0.796	0.015
LDL	2000	0.972	0.787	0.015
Direct	500	0.450	0.792	0.075
Direct	1000	0.500	0.787	0.140
Direct	2000	0.550	0.770	0.145

TABLE C4. Metrics for Werner-like — Local Broadcasting.

Model	Train Size	Accuracy	Fidelity	FID
Cholesky	500	0.972	0.962	0.004
Cholesky	1000	0.926	0.961	0.004
Cholesky	2000	0.976	0.962	0.003
LDL	500	0.990	0.959	0.004
LDL	1000	0.945	0.959	0.005
LDL	2000	0.975	0.962	0.003
Direct	500	0.300	0.963	0.047
Direct	1000	0.320	0.962	0.041
Direct	2000	0.350	0.962	0.055

sented in Fig. C1, we performed an asymptotic analysis of the computational costs. This analysis validates the empirically observed $\Theta(d^2)$ and $\Theta(d^3)$ scaling behaviors and clarifies the trade-offs between the direct and decomposition-based approaches when extending the

TABLE C5. Metrics for Werner-like — Nonlocal Broadcasting.

Model	Train Size	Accuracy	Fidelity	FID
Cholesky	500	0.945	0.943	0.014
Cholesky	1000	0.974	0.941	0.008
Cholesky	2000	0.977	0.940	0.008
LDL	500	0.981	0.940	0.010
LDL	1000	0.983	0.941	0.008
LDL	2000	0.983	0.941	0.010
Direct	500	0.400	0.940	0.068
Direct	1000	0.450	0.940	0.055
Direct	2000	0.500	0.941	0.064

TABLE C6. Metrics for Werner-like — Teleportable.

Model	Train Size	Accuracy	Fidelity	FID
Cholesky	500	0.968	0.922	0.017
Cholesky	1000	0.963	0.907	0.008
Cholesky	2000	1.000	0.926	0.138
LDL	500	0.979	0.922	0.015
LDL	1000	0.959	0.907	0.010
LDL	2000	0.990	0.926	0.014
Direct	500	0.500	0.926	0.050
Direct	1000	0.550	0.908	0.046
Direct	2000	0.600	0.927	0.051

framework to higher Hilbert space dimensions d .

We analyze the computational cost of generating K density matrices $\rho \in \mathbb{C}^{d \times d}$ during inference, where d is the Hilbert space dimension ($d = 4$ for the two-qubit case considered in the main text). Each generator is implemented as a multilayer perceptron (MLP) with parameter count

$$W(d) = \sum_{l=1}^L (n_{l-1}n_l + n_l).$$

a. Direct Generator. The direct model outputs $2d^2$ real parameters (corresponding to the real and imaginary entries of ρ). Generation cost is therefore dominated by the MLP forward pass:

$$\mathcal{T}_{\text{Direct}}(K) = \Theta(KW(d)).$$

Because the final affine map must produce $2d^2$ outputs,

$$W(d) \geq \Omega(d^2).$$

Thus, $\mathcal{T}_{\text{Direct}}(K) = \Omega(Kd^2)$ always holds. If hidden widths scale sub-cubically, $W(d) = o(d^3)$, the direct model avoids any $\Omega(d^3)$ term at inference.

b. Decomposition Generators. The decomposition models output $O(d^2)$ parameters that define a factorization (e.g., Cholesky or LDL). For each state, the following steps are required:

1. MLP forward pass: $\Theta(W(d))$,
2. One dense $d \times d$ complex matrix multiplication to

form $\rho: \Theta(d^3)$,
 3. Trace normalization and any additional assembly: $o(d^3)$.

The per-state complexity is therefore

$$\mathcal{T}_{\text{Decomp}}(K) = \Theta(K(W(d) + d^3)).$$

In particular, $\mathcal{T}_{\text{Decomp}}(K) \geq \Omega(Kd^3)$, as at least one $d \times d$ multiplication is unavoidable.

c. Implications. Both approaches scale linearly in K , but their dependence on d differs:

1. Direct: $\Theta(KW(d))$ with $W(d) \geq \Omega(d^2)$, which is asymptotically more efficient in d whenever $W(d) = o(d^3)$.
2. Decomposition: $\Theta(K(W(d) + d^3))$, with the d^3 term dominating unless the network width already grows at least cubically.

As a concrete example, at $d = 8$ one dense multiplication requires $d^3 = 512$ complex multiplications, or approximately 4096 real floating-point operations. Decomposition models guarantee physicality by construction but incur this $\Theta(d^3)$ overhead, whereas the direct model trades physical guarantees for potentially lower inference cost.

d. Empirical validation. To corroborate these asymptotic estimates, we performed machine-agnostic micro-benchmarks using synthetic generators without training. We measured per-state wall-clock time for forward passes, decomposition assembly, and validation (PSD and PPT checks) across a broad range of Hilbert space dimensions $d \in \{4, 6, 8, 12, \dots, 256\}$ and batch size $B = 1024$, on both CPU and GPU backends. Figure C1 summarizes the results. Forward-pass costs (Direct, Cholesky, LDL) scale nearly quadratically with d , consistent with the $\Theta(d^2)$ dependence on parameter count. By contrast, decomposition assembly (LL^\dagger , LDL^\dagger) and eigenvalue-based PSD/PPT checks exhibit cubic growth $\Theta(d^3)$, as predicted by the analytic model. The same trends hold on CPU and GPU, confirming that the complexity analysis is robust and machine-agnostic. Error bars denote 95% confidence intervals across 10 repeats.

e. Benchmark configuration. All scalability measurements were performed on both CPU and GPU backends using synthetic generators without training. We evaluated Hilbert space dimensions $d \in \{4, 8, 12, 16, 20, 24, 28, 32, 36, 40, 64, 96, 128, 160, 256, 512\}$ and batch sizes $B \in \{64, 128, 256, 512, 1024\}$. Each measurement was averaged over 15 independent repeats, with 95% confidence intervals reported as error bars. CPU experiments were restricted to four threads for consistency, and PPT checks were omitted beyond $d = 32$ to prevent prohibitive cubic overhead. Since models were not trained, these benchmarks isolate the inference and validation cost of forward passes, decomposition assembly, and Positive Semi Definite checks.

Appendix D: Mathematical Derivations

1. Equivalence of Entanglement and Teleportation fidelity for Bell-Diagonal States

Definition. Bell-diagonal states have the density matrix form:

$$\rho_{\text{Bell}} = \frac{1}{4} \left(\mathbb{I} \otimes \mathbb{I} + \sum_{i=1}^3 c_i \sigma_i \otimes \sigma_i \right), \quad (\text{D1})$$

with parameters satisfying $-1 \leq c_i \leq 1$, and σ_i being the Pauli matrices.

Entanglement criterion (Peres-Horodecki). A Bell-diagonal state is entangled if and only if its partial transpose has at least one negative eigenvalue. For this class of states, this is equivalent to:

$$|c_1| + |c_2| + |c_3| > 1. \quad (\text{D2})$$

Teleportation fidelity criterion (Horodecki). The maximal teleportation fidelity $F_{\text{max}}(\rho)$ for Bell-diagonal states is given by:

$$F_{\text{max}}(\rho) = \frac{1}{2} \left(1 + \frac{|c_1| + |c_2| + |c_3|}{3} \right). \quad (\text{D3})$$

A quantum advantage (i.e., the state is a useful resource for teleportation) is achieved when $F_{\text{max}}(\rho) > \frac{2}{3}$.

Proof of equivalence. Substituting the fidelity condition and simplifying, we obtain:

$$\frac{1}{2} \left(1 + \frac{|c_1| + |c_2| + |c_3|}{3} \right) > \frac{2}{3} \quad (\text{D4})$$

$$1 + \frac{|c_1| + |c_2| + |c_3|}{3} > \frac{4}{3} \quad (\text{D5})$$

$$\frac{|c_1| + |c_2| + |c_3|}{3} > \frac{1}{3} \quad (\text{D6})$$

$$|c_1| + |c_2| + |c_3| > 1. \quad (\text{D7})$$

2. Equivalence of Teleportation fidelity and PPT Inseparability for Werner-like Two-Qubit States

For the Werner-like family of two-qubit states considered in this work, the condition for teleportation fidelity coincides exactly with the PPT inseparability condition.

a. State definition. Consider the two-qubit state

$$\rho(p, \alpha) = p |\psi\rangle\langle\psi| + \frac{1-p}{4} I_4, \quad 0 \leq p \leq 1, \quad (\text{D8})$$

where

$$|\psi\rangle = \alpha|00\rangle + \beta|11\rangle, \quad \alpha, \beta \geq 0, \quad \alpha^2 + \beta^2 = 1. \quad (\text{D9})$$

In Bloch form, the diagonal correlation tensor is

$$T = \text{diag}(2p\alpha\beta, -2p\alpha\beta, p). \quad (\text{D10})$$




FIG. C1. Empirical scaling of per-state inference and validation costs with Hilbert dimension d at batch size $B = 1024$. Top row: CPU results. Bottom row: GPU results. Left: Forward-pass costs for Direct, Cholesky, and LDL generators, scaling $\sim d^2$. Right: Assembly (LL^\dagger , LDL^\dagger) and PSD/PPT eigenvalue checks, scaling $\sim d^3$. Error bars indicate 95% confidence intervals across 10 repeats.

b. Teleportation Fidelity. For a two-qubit state with diagonal correlation tensor T , the maximal teleportation fidelity is

$$F_{\max}(\rho) = \frac{1}{2} \left(1 + \frac{1}{3} N(\rho) \right), \quad (\text{D11})$$

$$N(\rho) = \sum_{i=1}^3 |t_i| = p(1 + 4\alpha\beta). \quad (\text{D12})$$

Teleportation is useful if $F_{\max} > 2/3$, equivalently $N(\rho) > 1$, yielding

$$p(1 + 4\alpha\beta) > 1. \quad (\text{D13})$$

c. PPT inseparability. The partial transpose ρ^{T_B} has smallest eigenvalue

$$\lambda_- = \frac{1-p}{4} - p\alpha\beta. \quad (\text{D14})$$

The state is entangled (PPT-negative) iff $\lambda_- < 0$, which reduces to

$$p(1 + 4\alpha\beta) > 1. \quad (\text{D15})$$

d. Equivalence. Comparing Eqs. (D13) and (D15), teleportation-useful and PPT-negative conditions coincide for this family. Both are governed by the boundary

$$p = \frac{1}{1 + 4\alpha\beta} = \frac{1}{1 + 4\sqrt{\alpha^2(1 - \alpha^2)}}. \quad (\text{D16})$$

The commonly cited threshold $p > 1/3$ corresponds to maximal entanglement ($\alpha^2 = \beta^2 = 1/2$) and is necessary but not sufficient when α is unrestricted. This equivalence is specific to the Werner-like two-qubit family.