Quantum Teleportation State Explorer
Full Theory, Algorithm, and User Manual

Updated for `Demo.Developing.14`

Abstract

These notes give a detailed theoretical, numerical, and practical explanation of the *Quantum Teleportation State Explorer*. The previous lecture-note theory has been preserved: Fock-basis parametrization, Wigner-function conventions, Gaussian teleportation channels, SLSQP optimization, constraints, cutoff diagnostics, seed states, robust objectives, and reference-state interpretation remain in the document. New sections have been added for the latest app version, `Demo.Developing.14`: computation profiles, live Wigner preview modes, comparison-sweep finesse, full-covariance Gaussian noise, correlated and rotated EPR imperfections, covariance-stability safeguards, export behavior, and a step-by-step user manual. The document should be read both as lecture notes and as a practical operating guide for the app.

Contents

1 How to Use the App: Practical Manual for `Demo.Developing.14`

This section is intentionally practical. The later sections give the underlying theory in detail, while this section explains how to operate the app and how to interpret each major page.

1.1 Basic workflow

A reliable run should usually follow this sequence:

1. Open the app and start from the **Welcome** page.
2. Go to **Settings / Run**.
3. Choose a **Computation profile**. The profiles are numerical templates only; they do not define the physics.
4. Press **Apply Profile** if you want the selected profile to fill runtime settings such as N_{cut}, grid size, number of starts, iteration limit, live-preview resolution, and comparison sweep points.
5. Manually configure the channel: squeezing r, gain, anisotropic noise, correlated EPR noise, rotated EPR noise, thermal noise, detector noise, loss, drift, and phase diffusion.
6. Manually configure constraints: fixed energy, parity, displacement, Fock support, photon statistics, quadrature moments, Wigner negativity, overlap bounds, or high-Fock tail controls.
7. Press **Start optimization**. The optimizer runs in a worker thread and sends live updates to the GUI.
8. Inspect the **Wigner functions**, **History and photon distribution**, and **Final results** tabs.
9. Press **Run Comparison** in the **Comparison & Validation** tab. The comparison evaluates the optimized state and reference states under the same configured channel.
10. Adjust **comparison sweep points** if the comparison curves look too coarse. More points give smoother plots but require more time.
11. Save the final state or comparison results only after inspection. The comparison section does not save automatically.

1.2 Computation profiles

The app no longer uses experimental presets. The physical parameters are manual. Computation profiles are one-shot numerical templates:

Profile	Main purpose	Typical effect
Fast Search	quick debugging	small grid, few starts, short run
Balanced	ordinary exploration	moderate cutoff and grid
Presentation	live demonstration	smoother live visuals
High Accuracy	final verification	larger grid, more starts, finer sweeps
Custom	user-defined settings	no automatic template assumption

After a profile is applied, the user should still be able to edit every physical and numerical parameter. The profile is not a constraint and should not lock entry boxes.

1.3 Live Wigner preview modes

The live-preview mode controls visualization cost:

Mode	Interpretation
Off	no live Wigner rendering during optimization
Final only	Wigner plots are produced at the end
Low-resolution	occasional cheaper Wigner previews
Full-resolution	expensive high-quality live Wigner previews

For fidelity-only objectives, live Wigner preview is diagnostic; it does not change the optimization. For Wigner-negativity constraints or Wigner-negativity survival objectives, Wigner functions are part of the objective itself and must be computed inside the optimization loop.

1.4 Main tabs

The **Settings / Run** tab contains all numerical, channel, and constraint controls. The **Wigner functions** tab displays input and teleported output Wigner functions. The **History and photon distribution** tab shows optimization history and photon-number probabilities. The **Time-dependent fields** tab interprets the optimized state as a single oscillator mode and plots dimensionless field quadratures. The **Final results** tab provides the full constraint report, optimized coefficients, reference-state comparison, and convergence diagnostics. The **Comparison & Validation** tab produces post-optimization comparison plots and exportable data.

1.5 Comparison sweep finesse

The comparison section contains a user-selectable sweep-point count. Increasing it makes curves such as

$$F(r), \qquad F(g), \qquad F(\eta, r), \qquad R_{\mathrm{surv}}(r) \tag{1}$$

more detailed. This changes only the comparison plots. It does not change the optimized state unless the optimization is run again.

1.6 Reading the final report

The final report should be read in this order. First check whether the state is feasible and whether the maximum constraint violation is small. Then compare the optimized fidelity with the coherent reference under the same channel. Then inspect tail probability and cutoff diagnostics. Finally, interpret Wigner negativity and negativity survival only if the state has meaningful input Wigner negativity. If the input negativity is zero, the survival ratio is undefined and should be reported as not applicable.

1.7 Reading comparison plots

The comparison plots answer different questions. Fidelity bars compare optimized and reference states at the configured channel. Wigner-negativity bars show whether nonclassicality survives teleportation. Survival-ratio plots are meaningful only for states with nonzero input negativity. Fidelity-versus-squeezing plots show how the same input state behaves as the finite-squeezing part of the channel is changed. Heat maps reveal robustness against pairs of noise parameters. Wigner-pair plots show visually how each input distribution is distorted by the channel.

1.8 Export behavior

The app separates calculation from export. Running comparison displays plots inside the GUI. It does not automatically create files. Pressing **Save Comparison Results** creates a timestamped folder containing plots, numerical tables, and metadata. Optimized-state export includes the Fock coefficients and enough channel/configuration information to reuse the state in external scripts.

2 The Problem in One Sentence

The computational problem is to find a quantum state $|\psi\rangle$ that maximizes its own teleportation fidelity through a chosen continuous-variable quantum channel:

$$\boxed{F(\psi) = \langle\psi|\ \mathcal{E}(|\psi\rangle\langle\psi|)\,|\psi\rangle\,.} \tag{2}$$

The state is not allowed to be arbitrary. It must be normalized and have a fixed mean photon number:

$$\boxed{\langle\psi|\psi\rangle = 1, \qquad \langle\psi|\,\hat{n}\,|\psi\rangle = N.} \tag{3}$$

The code then optionally adds further constraints, but the core mathematical problem is the fixed-energy constrained maximization of $F(\psi)$.

Remark 2.1. *The most important distinction is between the* physical problem *and the* finite-dimensional numerical problem. *The physical Hilbert space is infinite-dimensional. The computer solves a truncated version with $n = 0, 1, \ldots, N_{\text{cut}} - 1$. Much of the diagnostic machinery exists to decide whether the truncated optimum is a trustworthy approximation to the infinite-dimensional problem.*

3 Continuous-Variable Quantum States

3.1 Single-mode bosonic Hilbert space

The optimizer treats the input as a single bosonic mode. The annihilation and creation operators obey

$$[\hat{a}, \hat{a}^\dagger] = 1. \tag{4}$$

The Fock states $|n\rangle$ satisfy

$$\hat{n}\,|n\rangle = n\,|n\rangle\,, \tag{5}$$

$$\hat{a}\,|n\rangle = \sqrt{n}\,|n-1\rangle\,, \tag{6}$$

$$\hat{a}^\dagger\,|n\rangle = \sqrt{n+1}\,|n+1\rangle\,, \tag{7}$$

where

$$\hat{n} = \hat{a}^\dagger\hat{a}. \tag{8}$$

A general pure state is

$$|\psi\rangle = \sum_{n=0}^{\infty} c_n\,|n\rangle\,. \tag{9}$$

The numerical version truncates this to

$$|\psi\rangle_{N_{\text{cut}}} = \sum_{n=0}^{N_{\text{cut}}-1} c_n\,|n\rangle\,. \tag{10}$$

3.2 Quadrature convention

The dimensionless quadratures are

$$\hat{x} = \frac{\hat{a} + \hat{a}^\dagger}{\sqrt{2}}, \qquad \hat{p} = \frac{\hat{a} - \hat{a}^\dagger}{i\sqrt{2}}. \tag{11}$$

They satisfy

$$[\hat{x}, \hat{p}] = i. \tag{12}$$

The vacuum state has

$$\Delta x^2 = \Delta p^2 = \frac{1}{2}. \tag{13}$$

The phase-space coordinate vector is

$$\boldsymbol{r} = \begin{pmatrix} x \\ p \end{pmatrix}. \tag{14}$$

3.3 Displacement and classical amplitude

The complex field amplitude is

$$\alpha_\psi = \langle \hat{a} \rangle_\psi. \tag{15}$$

In the Fock basis,

$$\boxed{\langle \hat{a} \rangle = \sum_{n=1}^{N_{\mathrm{cut}}-1} \sqrt{n}\, c_{n-1}^* c_n.} \tag{16}$$

The quadrature means are

$$\langle x \rangle = \sqrt{2}\, \mathrm{Re}\langle a \rangle, \qquad \langle p \rangle = \sqrt{2}\, \mathrm{Im}\langle a \rangle. \tag{17}$$

The quantity

$$N_{\mathrm{disp}} = |\langle a \rangle|^2 \tag{18}$$

measures the energy stored in coherent displacement. The remaining part

$$N_{\mathrm{fluct}} = \langle n \rangle - |\langle a \rangle|^2 \tag{19}$$

is a useful measure of nonclassical or fluctuation energy.

4 Wigner Functions

4.1 Definition and interpretation

The Wigner function is a real phase-space representation of a quantum state. It is not always nonnegative, so it is not a classical probability density. However, it reproduces quadrature marginal distributions and is extremely convenient for Gaussian channels.

In the convention used by the code, the vacuum Wigner function is

$$W_0(x, p) = \frac{1}{\pi} \exp(-x^2 - p^2). \tag{20}$$

The normalization convention is

$$\boxed{\int_{\mathbb{R}^2} W_\rho(x, p)\, \mathrm{d}x\, \mathrm{d}p = 1.} \tag{21}$$

4.2 Overlap formula

For two density operators ρ and σ,

$$\boxed{\mathrm{Tr}(\rho\sigma) = 2\pi \int_{\mathbb{R}^2} W_\rho(x,p) W_\sigma(x,p)\, \mathrm{d}x\, \mathrm{d}p.} \tag{22}$$

This identity is the reason the code can evaluate fidelity as an integral over two Wigner functions.

If the input is pure, $\rho_\psi = |\psi\rangle\langle\psi|$, and the output is

$$\rho_{\mathrm{out}} = \mathcal{E}(\rho_\psi), \tag{23}$$

then

$$F(\psi) = \langle\psi|\, \rho_{\mathrm{out}}\, |\psi\rangle = \mathrm{Tr}(\rho_\psi \rho_{\mathrm{out}}). \tag{24}$$

Therefore

$$\boxed{F(\psi) = 2\pi \int W_\psi(x,p) W_{\mathrm{out}}(x,p)\, \mathrm{d}x\, \mathrm{d}p.} \tag{25}$$

The numerical code approximates this by a rectangular-grid sum:

$$F(\psi) \approx 2\pi \sum_{i,j} W_\psi(x_i,p_j) W_{\mathrm{out}}(x_i,p_j) \Delta x \Delta p. \tag{26}$$

4.3 Wigner negativity

A common measure of non-Gaussianity is the Wigner negativity

$$\mathcal{N}_W = \frac{1}{2} \int \left(|W(x,p)| - W(x,p) \right)\, \mathrm{d}x\, \mathrm{d}p. \tag{27}$$

If $W \geq 0$ everywhere, then $\mathcal{N}_W = 0$. Many nonclassical states, such as Fock states and cat states, have $\mathcal{N}_W > 0$. The app can compute this quantity and can optionally constrain it, although such constraints are computationally expensive because they require repeated Wigner evaluations inside the optimizer.

5 Continuous-Variable Teleportation as a Channel

5.1 Teleportation in phase space

The implemented channel is a Gaussian teleportation-like channel of the form

$$\boldsymbol{r}_{\mathrm{out}} = g\boldsymbol{r}_{\mathrm{in}} + \boldsymbol{\xi}, \qquad \boldsymbol{\xi} = \begin{pmatrix} \xi_x \\ \xi_p \end{pmatrix}, \tag{28}$$

where g is the classical gain and $\boldsymbol{\xi}$ is Gaussian noise. The original diagonal-noise model used

$$Y = \begin{pmatrix} \nu_x & 0 \\ 0 & \nu_p \end{pmatrix}, \tag{29}$$

while the extended app permits the full covariance matrix

$$\boxed{Y = \begin{pmatrix} \nu_x & \nu_{xp} \\ \nu_{xp} & \nu_p \end{pmatrix}.} \tag{30}$$

Thus,

$$\mathbb{E}\xi_x = \mathbb{E}\xi_p = 0, \qquad \mathbb{E}\xi_x^2 = \nu_x, \qquad \mathbb{E}\xi_p^2 = \nu_p, \qquad \mathbb{E}\xi_x\xi_p = \nu_{xp}. \tag{31}$$

The off-diagonal term ν_{xp} models correlated imperfections in the EPR resource, rotated EPR-noise axes, or correlated classical feed-forward errors.

5.2 Gain transformation

If the channel only applied gain, then the Wigner function would transform as

$$W_g(x,p) = \frac{1}{g^2} W_{\text{in}}\left(\frac{x}{g}, \frac{p}{g}\right). \tag{32}$$

The factor $1/g^2$ preserves normalization:

$$\int W_g(x,p)\,\mathrm{d}x\,\mathrm{d}p = 1. \tag{33}$$

In the code, this is evaluated by interpolation on the phase-space grid.

5.3 Additive Gaussian noise

The noise kernel is

$$G_{\nu_x,\nu_p}(x,p) = \frac{1}{2\pi\sqrt{\nu_x\nu_p}} \exp\left[-\frac{1}{2}\left(\frac{x^2}{\nu_x} + \frac{p^2}{\nu_p}\right)\right]. \tag{34}$$

The output Wigner function is

$$W_{\text{out}}(x,p) = \int G_{\nu_x,\nu_p}(x-x', p-p') W_g(x',p')\,\mathrm{d}x'\,\mathrm{d}p'. \tag{35}$$

This convolution is evaluated with FFT convolution. The discrete formula is

$$W_{\text{out}} \approx \text{fftconvolve}(W_g, G, \texttt{mode=same})\Delta x \Delta p. \tag{36}$$

For the full covariance model,

$$Y = \begin{pmatrix} \nu_x & \nu_{xp} \\ \nu_{xp} & \nu_p \end{pmatrix}, \qquad D = \det Y = \nu_x\nu_p - \nu_{xp}^2, \tag{37}$$

the Gaussian kernel becomes

$$\boxed{G_Y(x,p) = \frac{1}{2\pi\sqrt{D}} \exp\left[-\frac{1}{2D}\left(\nu_p x^2 - 2\nu_{xp}xp + \nu_x p^2\right)\right].} \tag{38}$$

The code uses this full kernel whenever correlated or rotated EPR noise is enabled. The diagonal kernel is recovered by setting $\nu_{xp} = 0$.

5.4 Finite-squeezing noise

For ideal Braunstein–Kimble teleportation at unit gain, the finite squeezing of the entangled resource introduces isotropic Gaussian noise

$$\boxed{\nu_x = \nu_p = e^{-2r}.} \tag{39}$$

Here r is the two-mode squeezing parameter. Larger r means better entanglement and less teleportation noise.

For a coherent state under unit-gain isotropic noise ν, the benchmark fidelity is

$$\boxed{F_{\text{coh}} = \frac{1}{1+\nu}.} \tag{40}$$

For $r = 0.9$,

$$\nu = e^{-1.8} \approx 0.1653, \qquad F_{\text{coh}} \approx 0.858. \tag{41}$$

This formula is an important sanity check. If the code is run with unit gain, isotropic finite-squeezing noise, and no additional constraints, the coherent state should be recovered or tied by an equivalent displaced vacuum.

5.5 Anisotropic noise

The channel may have

$$\nu_x \neq \nu_p. \tag{42}$$

This is physically important because the optimal state may then redistribute its uncertainty between x and p. A squeezed or non-Gaussian state can outperform a coherent state for anisotropic channels, because coherent states have fixed quadrature variances

$$\Delta x^2 = \Delta p^2 = \frac{1}{2}. \tag{43}$$

5.6 Correlated non-ideal EPR noise

A non-ideal EPR resource need not add independent noise to x and p. Imperfect mode matching, correlated electronics, or an EPR covariance ellipse that is not aligned with the measurement axes can produce

$$\nu_{xp} \neq 0. \tag{44}$$

The app implements the correlated model through a correlation coefficient ρ:

$$\boxed{\nu_{xp} = \rho\sqrt{\nu_x \nu_p}}, \qquad -1 < \rho < 1. \tag{45}$$

The effective covariance is therefore

$$Y = \begin{pmatrix} \nu_x & \rho\sqrt{\nu_x \nu_p} \\ \rho\sqrt{\nu_x \nu_p} & \nu_p \end{pmatrix}. \tag{46}$$

When $\rho > 0$, positive x-noise tends to appear with positive p-noise; when $\rho < 0$, the two noise quadratures tend to fluctuate in opposite directions. In Wigner space this shears and tilts the Gaussian convolution kernel.

5.7 Rotated asymmetric EPR noise

A more general EPR-resource imperfection is a rotated noise ellipse. Let the principal noise variances be

$$\nu_+ \quad \text{and} \quad \nu_-, \tag{47}$$

and let the principal axes be rotated by an angle θ relative to the x, p axes. With

$$R(\theta) = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}, \tag{48}$$

the added rotated covariance is

$$\boxed{Y_{\text{rot}} = R(\theta) \begin{pmatrix} \nu_+ & 0 \\ 0 & \nu_- \end{pmatrix} R(\theta)^T.} \tag{49}$$

Explicitly,

$$\nu_x^{(\text{rot})} = \nu_+ \cos^2\theta + \nu_- \sin^2\theta, \tag{50}$$

$$\nu_p^{(\text{rot})} = \nu_+ \sin^2\theta + \nu_- \cos^2\theta, \tag{51}$$

$$\nu_{xp}^{(\text{rot})} = (\nu_+ - \nu_-) \sin\theta \cos\theta. \tag{52}$$

This model includes both anisotropy and correlation. It is useful when the noisy EPR squeezing and anti-squeezing axes are not aligned with the homodyne or feed-forward coordinates.

5.8 Loss, mode mismatch, and environmental noise

Optical loss or mode mismatch can be represented as a beam-splitter interaction with an environmental mode. If T is the transmissivity, then

$$\boldsymbol{r}_{\text{out}} = \sqrt{T}\,\boldsymbol{r}_{\text{in}} + \sqrt{1-T}\,\boldsymbol{r}_{\text{env}}. \tag{53}$$

For a thermal environment with mean occupation $\bar{n}_{\text{env}}$,

$$\Delta x^2_{\text{env}} = \Delta p^2_{\text{env}} = \frac{2\bar{n}_{\text{env}}+1}{2}. \tag{54}$$

The app folds this into the teleportation channel by replacing

$$g \mapsto \sqrt{T}\,g, \tag{55}$$

$$\nu_x \mapsto \nu_x + (1-T)\frac{2\bar{n}_{\text{env}}+1}{2}, \tag{56}$$

$$\nu_p \mapsto \nu_p + (1-T)\frac{2\bar{n}_{\text{env}}+1}{2}. \tag{57}$$

This is a physically motivated alternative to simply adding arbitrary isotropic noise.

5.9 Deterministic displacement drift

Calibration errors in classical feed-forward can also produce a deterministic phase-space displacement. The model is

$$\boldsymbol{r}_{\text{out}} = g\boldsymbol{r}_{\text{in}} + \boldsymbol{\xi} + \boldsymbol{d}, \qquad \boldsymbol{d} = \begin{pmatrix} d_x \\ d_p \end{pmatrix}. \tag{58}$$

In Wigner space this is a translation:

$$W_{\text{out}}(x,p) = W_{\text{pre}}(x-d_x, p-d_p). \tag{59}$$

Unlike Gaussian jitter, this is not random broadening. It is a systematic offset, so it can reduce fidelity even when the noise covariance is small.

5.10 Thermal and detector noise

The code includes phenomenological noise additions:

$$\nu_T = s_T(2\bar{n}_T + 1), \tag{60}$$

where $\bar{n}_T$ is a thermal occupation parameter and s_T is a strength coefficient. Detector inefficiency is modeled by

$$\nu_D = s_D\frac{1-\eta}{\eta}, \tag{61}$$

where η is detector efficiency and s_D is a user-chosen strength. These are not meant to be the only possible detector models; they are simple additive-noise controls for exploring imperfections.

5.11 Complete positivity

A one-mode Gaussian channel has the covariance transformation

$$V \mapsto XVX^T + Y. \tag{62}$$

For the present channel,

$$X = gI, \qquad Y = \begin{pmatrix} \nu_x & \nu_{xp} \\ \nu_{xp} & \nu_p \end{pmatrix}. \tag{63}$$

A necessary and sufficient complete-positivity condition is

$$Y + i\Omega - iX\Omega X^T \geq 0, \qquad \Omega = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \tag{64}$$

For $X = gI$ and a full real covariance matrix Y, this reduces to

$$\boxed{\det Y = \nu_x \nu_p - \nu_{xp}^2 \geq \frac{(1-g^2)^2}{4}.} \tag{65}$$

The diagonal formula

$$\nu_x \nu_p \geq \frac{(1-g^2)^2}{4} \tag{66}$$

is recovered by setting $\nu_{xp} = 0$. The app optionally enforces the full condition by increasing the diagonal noise if the user-selected covariance would otherwise be unphysical.

5.12 Phase diffusion

Phase diffusion is the channel

$$\rho \mapsto \int P_\sigma(\phi) e^{-i\phi\hat{n}} \rho e^{i\phi\hat{n}} \, \mathrm{d}\phi, \tag{67}$$

where

$$P_\sigma(\phi) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\phi^2}{2\sigma^2}\right). \tag{68}$$

In Wigner space, $e^{-i\phi\hat{n}}$ rotates the phase-space distribution. Therefore phase diffusion is an angular average:

$$W_{\text{diff}}(x,p) = \int P_\sigma(\phi) W(R_{-\phi}(x,p)) \, \mathrm{d}\phi. \tag{69}$$

The code computes this integral by Gauss–Hermite quadrature.

6 The Variational Optimization Problem

6.1 Finite-dimensional problem

After truncation, the state vector is

$$\boldsymbol{c} = (c_0, c_1, \ldots, c_{N_{\text{cut}}-1})^T \in \mathbb{C}^{N_{\text{cut}}}. \tag{70}$$

The feasible set is

$$\mathcal{M}_{N,N_{\text{cut}}} = \left\{ \boldsymbol{c} \in \mathbb{C}^{N_{\text{cut}}} : \sum_n |c_n|^2 = 1, \quad \sum_n n|c_n|^2 = N \right\}. \tag{71}$$

The finite-dimensional optimization problem is

$$\boxed{\max_{\boldsymbol{c} \in \mathcal{M}_{N,N_{\text{cut}}}} F(\boldsymbol{c}).} \tag{72}$$

6.2 Why optimize coefficients instead of Wigner functions?

One could imagine optimizing over $W(x,p)$ directly. This is dangerous because most real functions are not valid Wigner functions of quantum states. They may violate positivity of the corresponding density matrix, uncertainty relations, or normalization. By optimizing over $|\psi\rangle$ directly, every trial state is automatically physical after normalization.

Thus, the algorithm separates the problem into two roles:

Task	Representation	Reason
State validity	Fock coefficients c_n	Automatically physical pure states
Channel propagation	Wigner grid $W(x,p)$	Gaussian noise is easy by convolution
Fidelity evaluation	Wigner overlap integral	Efficient numerical integral

6.3 Real parameter vector

SLSQP is a real-variable optimizer. The complex coefficients are converted to a real vector

$$\boldsymbol{y} = (\operatorname{Re} c_0, \ldots, \operatorname{Re} c_{d-1}, \operatorname{Im} c_0, \ldots, \operatorname{Im} c_{d-1}) \in \mathbb{R}^{2d}, \tag{73}$$

where d is the number of active Fock states.

The objective function passed to SLSQP is

$$\boxed{f(\boldsymbol{y}) = -F(\boldsymbol{y}).} \tag{74}$$

Minimizing f is equivalent to maximizing fidelity.

6.4 How SLSQP works

It is useful to picture the optimization as motion in the real coefficient space

$$\boldsymbol{y} \in \mathbb{R}^{2d}. \tag{75}$$

For a state using d active Fock basis states, this is a $2d$-dimensional real space: the first d coordinates are the real parts of the amplitudes and the next d coordinates are the imaginary parts. The optimizer searches this space for a point that gives large teleportation fidelity while remaining on the constraint surface defined by normalization, fixed energy, and any optional constraints.

A simple unconstrained gradient descent/ascent method would update the coefficients approximately as

$$\boldsymbol{y}_{k+1} = \boldsymbol{y}_k + \eta \nabla F(\boldsymbol{y}_k), \tag{76}$$

where η is a step size. This picture is helpful intuitively, but it is incomplete for the app because the state cannot move freely in all directions. The allowed state must remain on the constrained manifold

$$g_i(\boldsymbol{y}) = 0, \qquad h_j(\boldsymbol{y}) \geq 0. \tag{77}$$

A naive gradient step usually breaks normalization and fixed energy. Therefore, the code uses SLSQP rather than ordinary gradient ascent.

SLSQP stands for *Sequential Least-Squares Quadratic Programming*. It is a local constrained optimization algorithm. Instead of taking a raw gradient step, it constructs a sequence of simpler quadratic programming problems. At iteration k, suppose the current coefficient vector is $\boldsymbol{y}_k$. The nonlinear objective is locally approximated by a quadratic model, and the nonlinear constraints are locally approximated by their first-order Taylor expansions. A typical SLSQP subproblem has the form

$$\min_{\boldsymbol{p}} \quad \nabla f(\boldsymbol{y}_k)^T \boldsymbol{p} + \frac{1}{2} \boldsymbol{p}^T B_k \boldsymbol{p}, \tag{78}$$

$$\text{subject to} \quad g_i(\boldsymbol{y}_k) + \nabla g_i(\boldsymbol{y}_k)^T \boldsymbol{p} = 0, \tag{79}$$

$$h_j(\boldsymbol{y}_k) + \nabla h_j(\boldsymbol{y}_k)^T \boldsymbol{p} \geq 0, \tag{80}$$

$$\boldsymbol{y}_{\min} - \boldsymbol{y}_k \leq \boldsymbol{p} \leq \boldsymbol{y}_{\max} - \boldsymbol{y}_k. \tag{81}$$

Here $\boldsymbol{p}$ is the proposed step direction and B_k is an approximation to the Hessian of the constrained Lagrangian. The equality constraints are linearized tangent-plane approximations to the exact constraint surface. The inequality constraints are also linearized, and only active or nearly active inequalities strongly influence the step.

After solving the quadratic subproblem, SLSQP does not necessarily take the full step. Instead, it chooses a step length α_k and updates

$$\boxed{\boldsymbol{y}_{k+1} = \boldsymbol{y}_k + \alpha_k \boldsymbol{p}_k.} \tag{82}$$

The step length is selected so that the objective and the constraint violations improve according to an internal merit function. In this sense, SLSQP behaves like a constrained gradient method: it tries to move toward larger fidelity, but the direction is corrected so that the state remains compatible with the constraints.

For the present problem, the main ingredients are:

(i) $f(\boldsymbol{y}) = -F(\boldsymbol{y})$, so decreasing f means increasing fidelity.

(ii) $g_1(\boldsymbol{y}) = \sum_n |c_n|^2 - 1$ enforces normalization.

(iii) $g_2(\boldsymbol{y}) = \sum_n n|c_n|^2 - N$ enforces fixed mean photon number.

(iv) Additional constraints, such as zero displacement or Wigner-negativity bounds, are added as extra g_i or h_j functions.

(v) Bounds $-2 \leq y_\nu \leq 2$ keep the real optimization variables in a numerically reasonable range. They do not determine the physics; the physics is mainly determined by normalization, energy, and channel fidelity.

The code does not supply closed-form analytic gradients of F. Therefore, the numerical optimizer estimates derivatives by finite differences. For example, the derivative in coordinate direction ν is approximated by comparing

$$f(\boldsymbol{y}_k + \epsilon \boldsymbol{e}_\nu) \quad \text{and} \quad f(\boldsymbol{y}_k), \tag{83}$$

where $\boldsymbol{e}_\nu$ is a unit coordinate vector and ϵ is a small finite-difference step. This is why one SLSQP iteration can require many Wigner-function evaluations: the optimizer must evaluate the objective several times to estimate local derivatives in the $2d$-dimensional coefficient space.

The quadratic model should not be interpreted as the true global shape of the fidelity landscape. It is only a local approximation. Since the fidelity landscape is nonconvex, SLSQP can converge to different local optima depending on the starting state. This is the reason for the multi-start strategy. The app runs SLSQP from several physically motivated seeds and then chooses the best feasible final state.

6.5 Geometric picture of the SLSQP step

The normalization constraint places the coefficient vector on a unit sphere in $\mathbb{R}^{2d}$:

$$\sum_n |c_n|^2 = 1. \tag{84}$$

The fixed-energy constraint cuts this sphere by another hypersurface:

$$\sum_n n|c_n|^2 = N. \tag{85}$$

Thus, the feasible states form a curved manifold inside the full coefficient space. At a current feasible point, the gradients of the constraints define normal directions to this manifold. SLSQP approximately removes the forbidden normal components of the fidelity-gradient direction and produces a step that lies close to the tangent space of the feasible set.

With only equality constraints, a useful linearized picture is

$$J_g(\boldsymbol{y}_k)\boldsymbol{p} = 0, \tag{86}$$

where J_g is the Jacobian matrix of the equality constraints. This equation means that the step $\boldsymbol{p}$ should not change normalization or energy to first order. The SLSQP subproblem chooses among such nearly tangent steps the one that most improves the local quadratic model of $-F$.

This geometric interpretation also explains why intermediate live states can sometimes look imperfect. The optimizer may temporarily evaluate nearby points that are not exactly feasible while estimating derivatives or searching for a step. The final result should be judged by the final constraint report, not by every sampled point in the live history.

6.6 Stopping criteria and interpretation of success

SLSQP stops when one of several numerical criteria is satisfied. In the app, the most important user-facing controls are the maximum number of iterations and the tolerance ftol. The tolerance affects how small the expected improvement, constraint residuals, and step changes must become before the solver considers the optimization converged.

A successful SLSQP message means that the solver found a local constrained optimum according to its finite-dimensional numerical criteria. It does not prove global optimality. The app therefore adds several safeguards:

(i) multiple initial seeds,

(ii) reference-state comparison,

(iii) explicit feasibility checks,

(iv) cutoff-tail diagnostics,

(v) optional local optimality probes,

(vi) optional cutoff scans over several values of N_{cut}.

Together these checks distinguish three different situations:

Observation	Meaning	Response
High fidelity and small constraint residuals	Valid finite-cutoff local optimum	Compare with references
Higher sampled point but constraint violation	Infeasible transient point	Do not select as final state
State changes strongly with N_{cut}	Possible cutoff-dependent sequence	Run convergence diagnostics

Thus, the best mental model is not "ordinary gradient ascent", but rather "local constrained quadratic search in the $2d$-dimensional coefficient space, repeated from many starting points."

6.7 Equality constraints

The always-active equality constraints are

$$g_1(\boldsymbol{y}) = \sum_n |c_n|^2 - 1 = 0, \tag{87}$$

$$g_2(\boldsymbol{y}) = \sum_n n|c_n|^2 - N = 0. \tag{88}$$

If zero displacement is enabled, the constraints

$$\mathrm{Re}\langle a\rangle = 0, \qquad \mathrm{Im}\langle a\rangle = 0 \tag{89}$$

are added.

6.8 Inequality constraints

Optional constraints are written in SLSQP form

$$h_j(\boldsymbol{y}) \geq 0. \tag{90}$$

For example, a high-Fock tail constraint

$$P_{\mathrm{tail}} \leq P_{\max} \tag{91}$$

is passed as

$$h(\boldsymbol{y}) = P_{\max} - P_{\mathrm{tail}}(\boldsymbol{y}) \geq 0. \tag{92}$$

A minimum Wigner negativity constraint

$$\mathcal{N}_W \geq \mathcal{N}_{\min} \tag{93}$$

is passed as

$$h(\boldsymbol{y}) = \mathcal{N}_W(\boldsymbol{y}) - \mathcal{N}_{\min} \geq 0. \tag{94}$$

6.9 KKT viewpoint

At a local optimum satisfying regularity conditions, the Karush–Kuhn–Tucker conditions hold. With equality constraints $g_i = 0$ and inequality constraints $h_j \geq 0$, the Lagrangian is

$$\mathcal{L}(\boldsymbol{y}, \lambda, \mu) = -F(\boldsymbol{y}) + \sum_i \lambda_i g_i(\boldsymbol{y}) - \sum_j \mu_j h_j(\boldsymbol{y}), \tag{95}$$

where

$$\mu_j \geq 0, \qquad \mu_j h_j(\boldsymbol{y}) = 0. \tag{96}$$

The stationarity condition is

$$\nabla_{\boldsymbol{y}} \mathcal{L} = 0. \tag{97}$$

The app does not symbolically solve these equations; SLSQP solves the nonlinear constrained problem numerically.

7 Euler–Lagrange and Functional Viewpoint

The original research question can be phrased as a functional optimization over Wigner functions:

$$F[W] = 2\pi \int W(x,p)(\mathcal{K}W)(x,p)\,\mathrm{d}x\,\mathrm{d}p, \tag{98}$$

where $\mathcal{K}$ represents the teleportation channel. Formally, an Euler–Lagrange equation could be written as

$$\frac{\delta}{\delta W}\left(F[W] - \lambda \int W - \mu E[W]\right) = 0. \tag{99}$$

However, this misses the hardest constraint: W must correspond to a positive semidefinite density operator. Therefore, direct Wigner-function variational calculus is not enough unless one also enforces quantum physicality.

The Fock-coefficient method is a finite-dimensional physical regularization of that functional problem. It automatically enforces positivity because

$$\rho = |\psi\rangle\langle\psi| \geq 0. \tag{100}$$

8 Seed States and Multi-Start Strategy

8.1 Need for multiple starts

The landscape $F(\boldsymbol{c})$ is nonconvex. Local solvers can become trapped in local maxima or stationary points. Therefore, the algorithm uses several starting states.

8.2 Coherent seed

A coherent state is

$$|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle\,. \tag{101}$$

In infinite dimension,

$$\langle n\rangle = |\alpha|^2. \tag{102}$$

In a finite cutoff, truncation changes the energy slightly. The code solves for α numerically so that

$$\langle\alpha|\hat{n}|\alpha\rangle_{N_{\mathrm{cut}}} = N. \tag{103}$$

8.3 Squeezed-vacuum seed

The squeezed vacuum is

$$|s\rangle = S(s)\,|0\rangle\,, \qquad S(s) = \exp\left[\frac{s}{2}(a^2 - a^{\dagger 2})\right]. \tag{104}$$

In infinite dimension,

$$\langle n\rangle = \sinh^2 s. \tag{105}$$

Again, the finite-cutoff code solves for s numerically.

8.4 Two-Fock seeds

A two-Fock seed is

$$|\psi_{m,n}\rangle = \sqrt{1-w}\,|m\rangle + e^{i\phi}\sqrt{w}\,|n\rangle\,. \tag{106}$$

It has energy

$$\langle n\rangle = (1-w)m + wn. \tag{107}$$

Therefore, to enforce energy N,

$$\boxed{w = \frac{N-m}{n-m}.} \tag{108}$$

This construction is useful because it gives exactly feasible seeds.

Remark 8.1. *A subtle bug can occur if integer target energy always returns the exact Fock state $|N\rangle$ for every random seed. Then the optimizer does not explore coherent-like states. The corrected algorithm includes coherent seeds and nontrivial two-Fock or random feasible seeds.*

9 Analytic Benchmark: Coherent State in the Isotropic Channel

Consider unit gain, isotropic additive noise $\nu_x = \nu_p = \nu$, and no phase diffusion. A coherent input has Wigner function

$$W_\alpha(x,p) = \frac{1}{\pi}\exp\left[-(x-x_0)^2 - (p-p_0)^2\right]. \tag{109}$$

The output is the same Gaussian convolved with a noise Gaussian, so its covariance increases. The overlap integral gives

$$F_{\text{coh}} = \frac{1}{1+\nu}. \tag{110}$$

For the finite-squeezing teleportation channel,

$$\nu = e^{-2r}. \tag{111}$$

Thus

$$F_{\text{coh}} = \frac{1}{1+e^{-2r}}. \tag{112}$$

This benchmark is independent of the coherent amplitude when $g = 1$ because the output displacement matches the input displacement.

9.1 Gain mismatch penalty

If $g \neq 1$, the output displacement is shifted from the input displacement. A coherent state with large $|\alpha|$ can then be strongly penalized because the output and input Gaussians are centered at different points. This is why non-displaced states may outperform coherent states in a gain-mismatched channel.

10 Constraints and Their Physical Meaning

10.1 Parity constraints

Even parity restricts

$$|\psi\rangle = \sum_{k=0}^{\infty} c_{2k}\,|2k\rangle\,, \tag{113}$$

while odd parity restricts

$$|\psi\rangle = \sum_{k=0}^{\infty} c_{2k+1} |2k+1\rangle . \tag{114}$$

Parity constraints are implemented by removing forbidden Fock indices from the active basis.

10.2 Photon statistics constraints

The photon-number variance is

$$\Delta n^2 = \langle n^2 \rangle - \langle n \rangle^2. \tag{115}$$

The Mandel parameter is

$$Q = \frac{\Delta n^2 - \langle n \rangle}{\langle n \rangle}. \tag{116}$$

The interpretations are

$$Q = 0 \quad \text{coherent-like Poissonian statistics,} \tag{117}$$

$$Q < 0 \quad \text{sub-Poissonian statistics,} \tag{118}$$

$$Q > 0 \quad \text{super-Poissonian statistics.} \tag{119}$$

10.3 Quadrature constraints

The second-order moments can be computed from

$$\langle x \rangle = \sqrt{2}\,\mathrm{Re}\langle a \rangle, \tag{120}$$

$$\langle p \rangle = \sqrt{2}\,\mathrm{Im}\langle a \rangle, \tag{121}$$

$$\langle a^2 \rangle = \sum_{n=2}^{N_{\mathrm{cut}}-1} \sqrt{n(n-1)}\, c_{n-2}^{*} c_n. \tag{122}$$

Then

$$\langle x^2 \rangle = \langle n \rangle + \frac{1}{2} + \mathrm{Re}\langle a^2 \rangle, \tag{123}$$

$$\langle p^2 \rangle = \langle n \rangle + \frac{1}{2} - \mathrm{Re}\langle a^2 \rangle, \tag{124}$$

$$\frac{1}{2}\langle xp + px \rangle = \mathrm{Im}\langle a^2 \rangle. \tag{125}$$

Thus

$$\Delta x^2 = \langle x^2 \rangle - \langle x \rangle^2, \tag{126}$$

$$\Delta p^2 = \langle p^2 \rangle - \langle p \rangle^2, \tag{127}$$

$$\mathrm{Cov}(x,p) = \frac{1}{2}\langle xp + px \rangle - \langle x \rangle \langle p \rangle. \tag{128}$$

10.4 Overlap constraints

To compare the optimized state with a reference state $|\phi\rangle$, the overlap is

$$O_\phi = |\, \langle \phi | \psi \rangle \,|^2. \tag{129}$$

A maximum coherent-overlap constraint deliberately excludes coherent-like states. A minimum coherent-overlap constraint deliberately searches within a coherent-like neighborhood. Such constraints are useful for controlled experiments, but they modify the mathematical problem.

11 Cutoff Effects and Infinite-Dimensional Subtleties

11.1 The cutoff is not physical

The Fock cutoff N_{cut} is a numerical approximation. A result is convincing only if the optimized state and fidelity stabilize as N_{cut} is increased.

11.2 Noncompactness under fixed mean energy

Fixed mean energy does not by itself prevent high-Fock excursions. Consider

$$|\psi_M\rangle = \sqrt{1 - \frac{N}{M}}\,|0\rangle + \sqrt{\frac{N}{M}}\,|M\rangle\,. \tag{130}$$

Then

$$\langle n\rangle = M\frac{N}{M} = N. \tag{131}$$

However, as $M \to \infty$, most of the probability returns to the vacuum while a tiny probability carries finite energy at very large photon number. Therefore, the infinite-dimensional fixed-energy pure-state problem may have maximizing sequences rather than stable maximizers.

11.3 Diagnostic quantities

The app monitors

$$P_{\mathrm{tail}} = \sum_{n=N_{\mathrm{cut}}-L}^{N_{\mathrm{cut}}-1} |c_n|^2, \tag{132}$$

$$E_{\mathrm{tail}} = \sum_{n=N_{\mathrm{cut}}-L}^{N_{\mathrm{cut}}-1} n|c_n|^2, \tag{133}$$

$$P_{\mathrm{upper}} = \sum_{n\geq 3N_{\mathrm{cut}}/4} |c_n|^2, \tag{134}$$

$$\mathrm{IPR}^{-1} = \frac{1}{\sum_n |c_n|^4}. \tag{135}$$

Large weight or energy near the cutoff is a warning that the optimizer may be exploiting the artificial boundary of the Hilbert space.

11.4 Feasible-best selection

During optimization, SLSQP may evaluate or pass through infeasible points. A live history plot can therefore show sampled states with higher objective values than the final feasible solution. The corrected algorithm distinguishes:

- sampled score: the objective value of whatever state was displayed;
- best feasible score: the best score among states satisfying all active constraints.

Only feasible states should be eligible for the final optimum.

12 Robust Objectives

Instead of optimizing fidelity for one channel parameter set, one can optimize an average over uncertain parameters:

$$\overline{F}(\psi) = \frac{1}{K}\sum_{k=1}^{K} F(\psi;\theta_k), \tag{136}$$

where θ_k may include different values of r, gain, noise, or phase diffusion.

A more conservative objective is the worst-case fidelity:

$$F_{\mathrm{worst}}(\psi) = \min_{1\le k\le K} F(\psi;\theta_k). \tag{137}$$

The first objective finds states that perform well on average. The second finds states with a better guaranteed performance under uncertainty.

13 Time-Dependent Field Interpretation

The optimized state can be visualized as a single oscillator mode. Define

$$X_\theta = \frac{ae^{-i\theta} + a^\dagger e^{i\theta}}{\sqrt{2}}. \tag{138}$$

With $\theta = \omega t$, the plotted field quadratures are

$$E(t) = E_0\langle X_{\omega t}\rangle, \tag{139}$$

$$B(t) = B_0\langle X_{\omega t+\pi/2}\rangle. \tag{140}$$

The uncertainty bands are

$$\Delta E(t) = E_0\sqrt{\langle X_{\omega t}^2\rangle - \langle X_{\omega t}\rangle^2}, \tag{141}$$

$$\Delta B(t) = B_0\sqrt{\langle X_{\omega t+\pi/2}^2\rangle - \langle X_{\omega t+\pi/2}\rangle^2}. \tag{142}$$

For coherent states, the mean fields oscillate sinusoidally. For number states or parity-symmetric states, the mean fields may vanish even though the fluctuations remain large.

14 Code Architecture

The app is divided into conceptual layers:

1. State layer: conversion between real optimizer variables and complex Fock coefficients.
2. Moment layer: energy, displacement, parity, photon statistics, quadrature moments.
3. Channel layer: gain, Gaussian convolution, phase diffusion.
4. Objective layer: Wigner fidelity and robust objective variants.
5. Constraint layer: normalization, energy, displacement, parity, tails, statistics, and overlaps.
6. Optimizer layer: multi-start SLSQP and feasible-best selection.
7. Diagnostic layer: cutoff support, convergence scans, reference-state comparison.
8. GUI layer: parameters, live plots, final reports, threading and queue messages.

Algorithm 1 Fidelity optimization algorithm

```
Read user-selected configuration.
Build active Fock basis.
Construct phase-space grid.
Build teleportation channel parameters.
Generate coherent, squeezed, Fock, two-Fock, and random seeds.
for each seed do
    Convert complex coefficients to real optimizer vector.
    Minimize −F using SLSQP subject to active constraints.
    Evaluate final state, Wigner functions, fidelity, and diagnostics.
    If state is feasible and improves the score, store it as best.
end for
Optionally polish the best state.
Compare against reference states.
Run convergence diagnostics.
Display final Wigner functions, photon distribution, history, and report.
```

15 Algorithm Summary

16 Common Failure Modes and Their Interpretation

16.1 The final state is not coherent in the default channel

Check whether any extra constraints or penalties are active. A tail constraint, bounded displacement constraint, phase diffusion, anisotropic noise, or gain mismatch changes the problem. The coherent state is the benchmark optimum only for the unconstrained, unit-gain, isotropic additive-noise channel.

16.2 The history plot reaches higher fidelity than the final state

This can happen if the high points are infeasible sampled states. The final answer must be selected from feasible states only. The report should show constraint feasibility and maximum violation.

16.3 Increasing N_{cut} changes the optimum

This may indicate either insufficient numerical convergence or a genuine cutoff-dependent maximizing sequence. Use cutoff diagnostics, scan N_{cut}, inspect tail probability, and compare reference fidelities.

16.4 The optimizer returns a Fock state unexpectedly

Check seed generation and constraints. If target energy is an integer and every seed collapses to $|N\rangle$, the optimizer may not explore coherent-like states. Include coherent and broad random feasible seeds.

17 Practical Recommended Workflows

17.1 Benchmark workflow

Use

$$g = 1, \qquad \nu_x = \nu_p = e^{-2r}, \qquad \sigma_{\mathrm{phase}} = 0, \tag{143}$$

with no extra constraints. The coherent state should be recovered.

17.2 Non-Gaussian discovery workflow

Enable anisotropic noise, phase diffusion, or gain mismatch. Run several N_{cut} values and compare the optimized state to coherent and squeezed references. Keep cutoff diagnostics active.

17.3 Convergence workflow

Run the same settings for

$$N_{\mathrm{cut}} = 8, 10, 12, 14, 16, 20, 24, 30. \tag{144}$$

Track

$$F_{\mathrm{opt}}, \quad P_{\mathrm{tail}}, \quad \Delta n^2, \quad \max\{n : P_n > 10^{-6}\}. \tag{145}$$

A stable optimum should not keep moving toward the cutoff.

18 New noise models in the presentation app

The presentation version contains the following channel imperfections beyond the original diagonal additive-noise model:

(i) correlated non-ideal EPR noise, controlled by an off-diagonal covariance ν_{xp} or equivalently by a correlation coefficient ρ;

(ii) rotated asymmetric EPR noise, represented by the covariance ellipse

$$Y_{\mathrm{rot}} = R(\theta)\,\mathrm{diag}(\nu_+, \nu_-)R(\theta)^T;$$

(iii) loss or mode mismatch to a vacuum/thermal environment, represented by transmissivity T and bath occupation $\bar{n}_{\mathrm{env}}$;

(iv) deterministic displacement drift, represented by a fixed translation (d_x, d_p) of the output Wigner function;

(v) the previous noise controls: finite squeezing, gain mismatch, diagonal anisotropic noise, thermal noise, detector noise, custom additive noise, and phase diffusion.

These models remain Gaussian at the channel level, so the channel is still evaluated efficiently as gain, full-covariance convolution, phase-diffusion averaging, and displacement on a Wigner grid.

19 Latest Channel and Numerical Updates in `Demo.Developing.14`

This section records the latest changes in the app without replacing the theory above. The earlier sections explain the mathematical objects; the present section states how the current implementation assembles and stabilizes them.

19.1 Finite squeezing as a baseline noise

In the current version, finite-squeezing noise is a baseline contribution. When finite squeezing is enabled,

$$\nu_{\mathrm{sq}} = e^{-2r} \tag{146}$$

is first added isotropically. Other noises are then added on top of it. In schematic form,

$$\begin{aligned} Y = \nu_{\mathrm{sq}} I + Y_{\mathrm{anis}} + Y_{\mathrm{thermal}} + Y_{\mathrm{detector}} + Y_{\mathrm{extra}} \\ + Y_{\mathrm{corr}} + Y_{\mathrm{rot}} + Y_{\mathrm{loss}}. \end{aligned} \tag{147}$$

Thus anisotropic noise does not replace finite squeezing. It modifies the already-present finite-squeezing channel. This is important for comparison sweeps: when the app plots fidelity versus r, changing r should change the finite-squeezing term even if other noise switches are active.

19.2 Full covariance and numerical stability

When correlated or rotated noise is enabled, the effective covariance matrix is

$$Y = \begin{pmatrix} \nu_x & \nu_{xp} \\ \nu_{xp} & \nu_p \end{pmatrix}. \tag{148}$$

The determinant

$$D = \nu_x \nu_p - \nu_{xp}^2 \tag{149}$$

must be positive for a well-conditioned Gaussian kernel. If D becomes too small, the convolution kernel becomes nearly line-like on the finite grid. This can create artificial jumps in comparison sweeps, especially at large r, where the finite-squeezing diagonal noise becomes small while a fixed off-diagonal covariance remains strong.

The current app therefore uses a covariance-stability guard. In practical terms, it prevents the effective correlation

$$\rho_{\text{eff}} = \frac{\nu_{xp}}{\sqrt{\nu_x \nu_p}} \tag{150}$$

from approaching exactly ± 1 and keeps D above a small numerical floor. This is a numerical regularization of the discretized Wigner-grid calculation, not a new physical noise source chosen by the user.

19.3 Complete positivity and stability are different checks

The complete-positivity condition

$$\det Y \geq \frac{(1-g^2)^2}{4} \tag{151}$$

ensures that the Gaussian map is a physical quantum channel. The covariance-stability guard is a separate finite-grid numerical check. A channel can satisfy complete positivity and still be numerically difficult if the covariance ellipse is extremely thin compared with the Wigner grid spacing.

19.4 Configured coherent benchmark versus ideal benchmark

The comparison section distinguishes two ideas. The ideal coherent benchmark

$$F_{\text{coh,ideal}}(r) = \frac{1}{1+e^{-2r}} \tag{152}$$

applies only to the unit-gain, isotropic, finite-squeezing-only channel. In a configured noisy channel, the coherent reference must be evaluated through the same channel as every other state. Therefore the comparison plot should not interpret the ideal curve as the actual benchmark when detector noise, thermal noise, correlated noise, rotated noise, phase diffusion, loss, or drift is enabled.

19.5 Why squeezing sweeps can jump

A sudden jump in a squeezing-dependent comparison curve usually indicates one of three issues:

(i) the covariance matrix became nearly singular at high r;

(ii) the Wigner grid was too coarse to resolve a narrow noise kernel;

(iii) the plotted ideal benchmark was being compared with a different configured channel.

The latest version addresses the first issue by stabilizing the covariance determinant, the second by allowing finer comparison sweeps and higher grid settings, and the third by separating configured-reference curves from ideal-reference curves.

19.6 Comparison sweep points

Let M be the selected comparison sweep-point number. A curve such as $F(r)$ is evaluated on M sample points between the chosen minimum and maximum values. Larger M gives smoother curves but costs more channel evaluations. For presentation plots, a moderate value such as 41 is usually visually smooth. For final checks, larger values may be useful.

19.7 State export and comparison export

The state export stores the optimized Fock coefficients, configuration, effective channel information, objective values, and metadata. The comparison export stores plot data, rendered PNGs, and JSON/CSV summaries in a timestamped folder. This separation is important: optimization produces a candidate state, while comparison tests how that state behaves against references and parameter sweeps.

20 Final Conceptual Summary

The optimizer is best understood as a variational quantum-state search. The state is represented in Hilbert space to guarantee physicality, while the channel and fidelity are evaluated in Wigner phase space because Gaussian teleportation noise is naturally represented there. The most important theoretical checks are the coherent-state benchmark for the unit-gain isotropic channel and the cutoff-convergence diagnostics for high-dimensional searches. When the optimizer finds non-Gaussian states, the key question is not whether they look coherent, but whether they are feasible, stable under cutoff increase, and better than relevant reference states under exactly the same channel and constraints.

A Appendix A: Useful Formulae

$$\langle a \rangle = \sum_{n=1}^{N_{\mathrm{cut}}-1} \sqrt{n} c_{n-1}^* c_n, \tag{153}$$

$$\langle a^2 \rangle = \sum_{n=2}^{N_{\mathrm{cut}}-1} \sqrt{n(n-1)} c_{n-2}^* c_n, \tag{154}$$

$$\langle n \rangle = \sum_n n |c_n|^2, \tag{155}$$

$$\langle n^2 \rangle = \sum_n n^2 |c_n|^2, \tag{156}$$

$$\Delta n^2 = \langle n^2 \rangle - \langle n \rangle^2, \tag{157}$$

$$Q = \frac{\Delta n^2 - \langle n \rangle}{\langle n \rangle}, \tag{158}$$

$$\langle x \rangle = \sqrt{2}\, \mathrm{Re} \langle a \rangle, \tag{159}$$

$$\langle p \rangle = \sqrt{2}\, \mathrm{Im} \langle a \rangle, \tag{160}$$

$$\Delta x^2 = \langle n \rangle + \frac{1}{2} + \mathrm{Re} \langle a^2 \rangle - \langle x \rangle^2, \tag{161}$$

$$\Delta p^2 = \langle n \rangle + \frac{1}{2} - \mathrm{Re}\langle a^2 \rangle - \langle p \rangle^2. \tag{162}$$

B Appendix B: Minimal Pseudocode

```
def objective(y):
    c = unpack_real_vector_to_fock_coefficients(y)
    c = normalize(c)
    W_in = qutip.wigner(ket2dm(c), xvec, pvec)
    W_out = apply_teleportation_channel(W_in, params)
    F = 2*pi*sum(W_in*W_out)*dx*dp
    return -F

constraints = [normalization_constraint, energy_constraint]
if zero_displacement:
    constraints += [real_a_constraint, imag_a_constraint]

for seed in seeds:
    result = scipy.optimize.minimize(
        objective,
        seed_to_real_vector(seed),
        method="SLSQP",
        constraints=constraints,
    )
    if result_is_feasible(result) and result.fidelity > best.fidelity:
        best = result
```