YQuantum: Alice and Bob

Haadi Khan, Arul Rhik Mazumder, Tanmay Neema, Ryan O'Farrell, Alice Wang

April 2025

1 Part 1

1.1 Part 1A

For Task 1a, we used QuTiP to create the Wigner functions and matplotlib to visualize them. Pictured here are the Fock states $|0\rangle, |1\rangle$ and $|2\rangle$, the coherent state $|\alpha\rangle$ with $\alpha = 2$, and the cat states $\frac{|\alpha\rangle+|-\alpha\rangle}{\sqrt{2}}$ and $\frac{|\alpha\rangle-|-\alpha\rangle}{\sqrt{2}}$. We also simulated the 3-cat state and dissipative cat state, pictures of which are available on the GitHub.

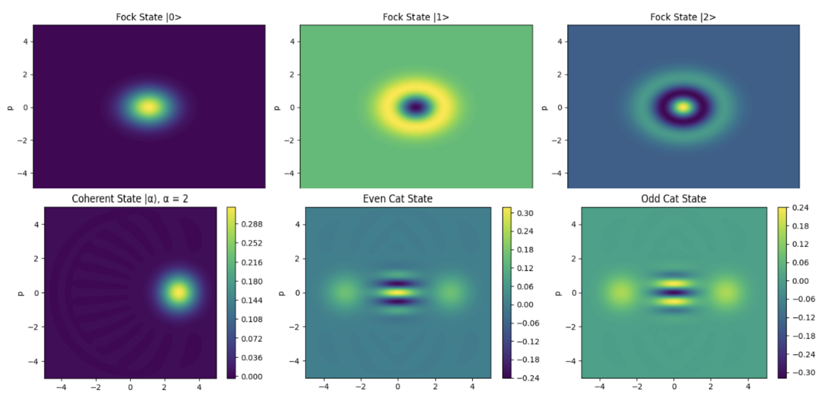

1.2 Part 1B

For Task 1b, we used the given least-sqauares approach to reconstruct the states from the Wigner functions. Our sampling algorithm detected the notable contours of the given Wigner function and sampled more heavily in regions where the Wigner function is far from zero.

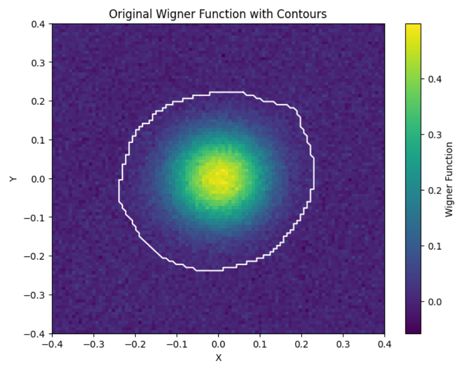

We then placed our sample points uniformly in both regions, with in total around 40 percent of the points concentrated inside the contour. Then, once we had our sample points, we could find the least squares fit given in the problem statement. We used SCS (the Splitting Cone Solver) to solve the convex optimization problem and find the optimal fit.

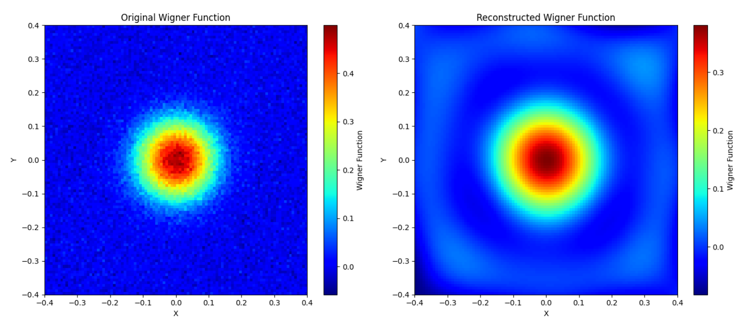

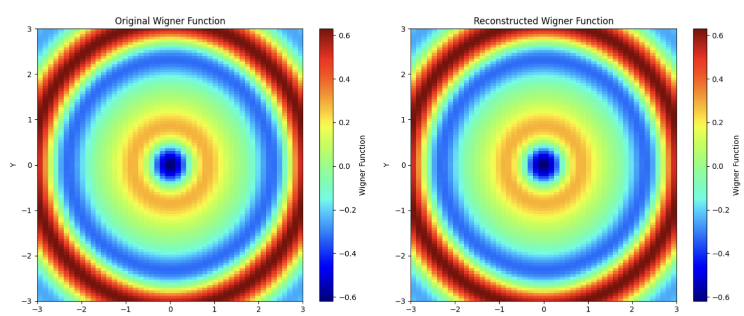

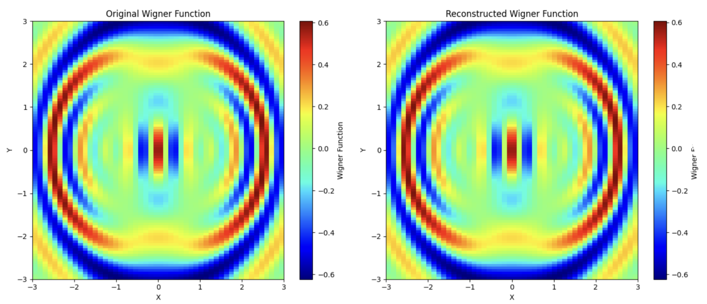

We did consider other reconstruction options. For example, we could use a Bayesian framework to quantify which density matrix is most likely to reproduce the observed measurements. With a uniform prior (using a uniform measure on all density matrices), we have for any given density matrix ρ that

$$P(\rho|w_1, \dots, w_N) \propto \prod_{i=1}^{N} P(w_i|\rho)$$

In particular, the posterior mode is exactly the maximum likelihood estimator, another possible reconstruction option. Assuming the measurements have a Gaussian error such that

$$W_\rho(a_k) - w_k \sim N(0, \sigma^2)$$

(which may or may not be accurate, it serves only as a calculation tool), we can compute this MLE to be

$$\arg\max_\rho e^{-\frac{1}{2\sigma^2}\sum_i (W_\rho(a_i) - w_i)^2}$$

which means that ρ is exactly the established least squares minimizer. Since this is independent of σ, we can take the $\sigma \to 0$ limit and recover the original case.

1.3 Part 1C

We first added Gaussian noise to our simulated Wigner functions and plotted the fidelity against the standard deviation of the noise. It is evident that noise quickly degrades the fidelity.

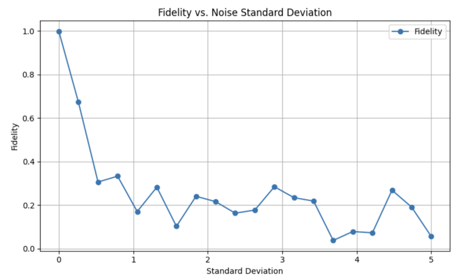

Figure 1: Fidelity vs. noise for a simulated coherent state

Then, we tested it on the real data. Observe that even on real Wigner functions the noiseless functions have a high fidelity, indicating that the least squares fit does indeed reproduce real-world states from Wigner functions.

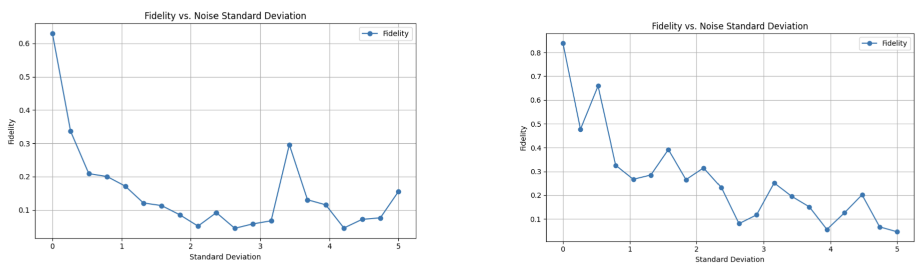

Figure 2: Real Fock (left) and cat (right) states

2 Part 2

2.1 Part 2A

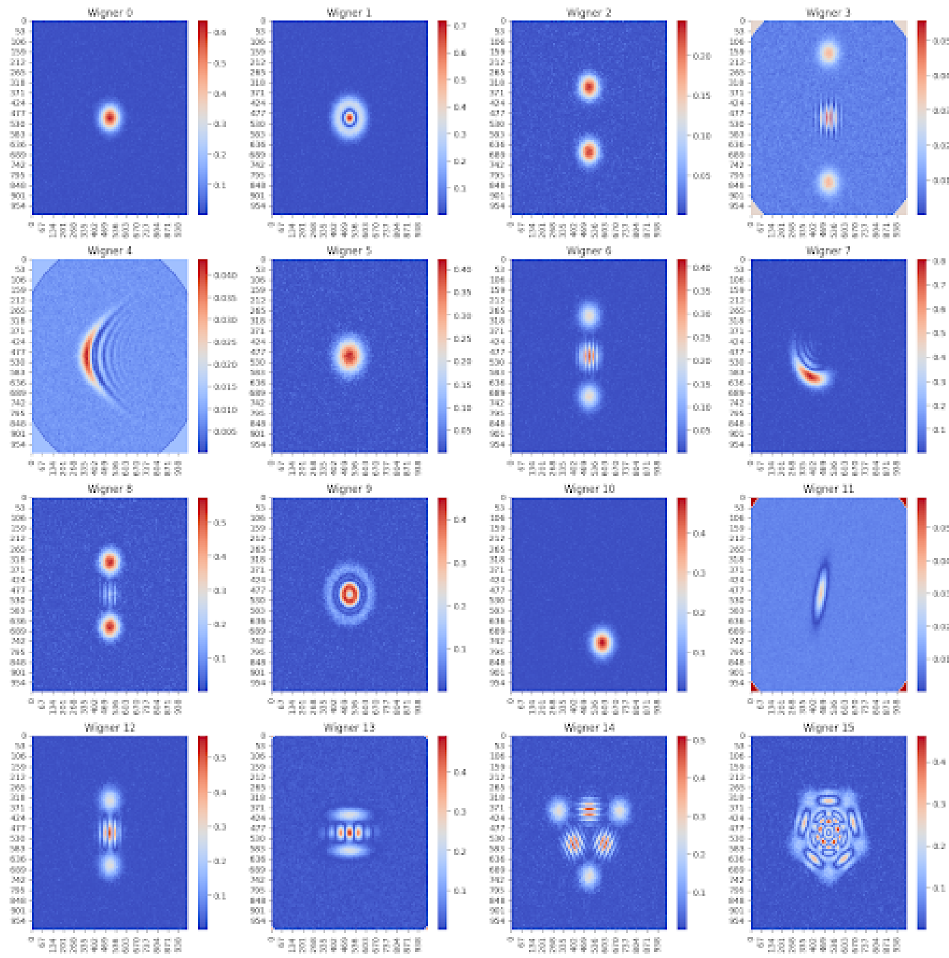

Figure 3: Wigner functions corrected for affine noise with Gaussian filtering.

Correcting affine transformations was straightforward via subtracting the average value of points along the edges of the phase space and normalizing the wavefunction. We applied Gaussian filtering with a heuristic of $\sigma = 3$, which was highly effective in correcting for everything except some shot noise (as shown in Figure 3). Our fidelity measurements will be included later.

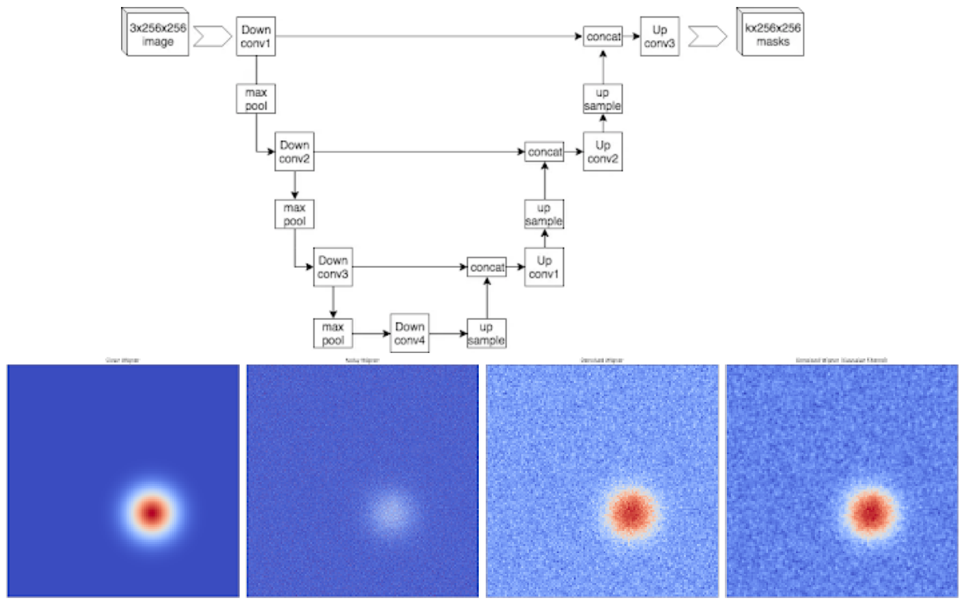

Figure 4: a) Top: representation of how U-Net works through multiple layers of ”downsampling” and ”upsampling”. b) Bottom: our AI’s denoising performance on a sample image in the test set.

2.2 Part 2B

To account for other types of noise, we generated a data set of roughly 3000 images. We analyzed 5 types of noise: Gaussian noise, Poisson noise, salt and pepper noise, speckle noise, and spatially-correlated noise (i.e. a moving average of Gaussian noise). For every combination of these types of noise, we generated 100 samples of random Wigner functions with this noise.

We used the convolutional neural network U-Net as our learning method. U-Net is a convolutional neural network designed for image segmentation with an encoder-decoder structure. The encoder path captures context through repeated convolution layers (”downsampling”), progressively reducing spatial dimensions while increasing feature depth. The decoder path compresses or ”upsamples” the feature maps to recover spatial resolution, combining them with corresponding high-resolution features from the encoder. This procedure, shown in Figure 4(a) makes U-Net highly effective for image segmentation even with limited data.

Our results for U-Net are demonstrated in Figure 4(b). We were able to successfully remove very large amounts of noise and reconstruct a relatively noise-less image. Our fidelity and loss measurements will be included later.

2.3 Part 2C

To improve performance (besides our original approach of sampling points within smaller, high-density regions), our team decided to use multithreading. However, the algorithm utilized in Part 1B is incompatible with multithreading. As a result, we used the CBC solver (instead of the SCS solver) for convex optimization. After some internal testing, we determined that 8 threads was best for our use case. We observed a 30x speedup across our reconstruction. We believe that a combination of the multithreading in addition to improvements of the CBC solver upon the SCS solver contributed to this improvement in performance.

References

[Agarwal, 1986] Agarwal, G. S. (1986). Wigner-function description of quantum noise ininterferometers. *Journal oF Modern Optics*.

[dynamiqs,] dynamiqs. Dynamiqs. http://www.dynamiqs.org/.

[Eddy Keming Chen, 2020] Eddy Keming Chen, R. T. (2020). Uniform probability distribution over all density matrices. *Quantum Studies: Mathematics and Foundations*.

[Hale and Townsend, 2014] Hale, N. and Townsend, A. (2014). Padua points in chebfun2. Last accessed 13 April 2025.

[qutip,] qutip. Qutip. http://www.qutip.org/.