

Generative Modeling of Quantum Distribution with Functional Flow Matching

Jachoon Hahm^{*12} Tak Hur^{*1} Joonseok Lee²³ Daniel K. Park¹⁴

Abstract

The emergence of powerful deep generative models based on diffusion and flow matching has enabled the learning and modeling of complex distributions. Learning quantum distributions, however, remains challenging due to the inherent difficulty of accurately modeling the meaningful physical properties of quantum states. We propose Quantum Flow Matching (QFM), a novel generative model designed to learn quantum distribution by utilizing spin Wigner function and flow matching. By converting density matrix into the spin Wigner function and leveraging functional flow matching to learn distributions in function space, QFM enables accurate and effective learning of multi-qubit quantum distributions. We demonstrate the effectiveness of our method by evaluating physical quantities such as trace, purity, and entanglement entropy of the generated quantum states, accurately capturing the underlying physics of the given quantum distributions.

1. Introduction

Despite the unprecedented success of deep generative models, such as diffusion models (Song et al., 2020; Ho et al., 2020) and flow matching (Lipman et al., 2022), learning the distributions of quantum states remains a challenging task. Leveraging state-of-the-art machine learning methods to model quantum states has been an interesting research direction (Carrasquilla et al., 2019; Carleo et al., 2019). However, applying modern generative models directly to learn quantum distributions has not been successful for several reasons. First, existing diffusion and flow matching methods are solely focused on learning representations of

^{*}Equal contribution ¹Department of Statistics and Data Science, Yonsei University, Seoul, Korea ²Graduate School of Data Science, Seoul National University, Seoul, Korea ³Google Research, Mountain View, California, United States ⁴Department of Applied Statistics, Yonsei University, Seoul, Korea. Correspondence to: Daniel K. Park <dkd.park@yonsei.ac.kr>.

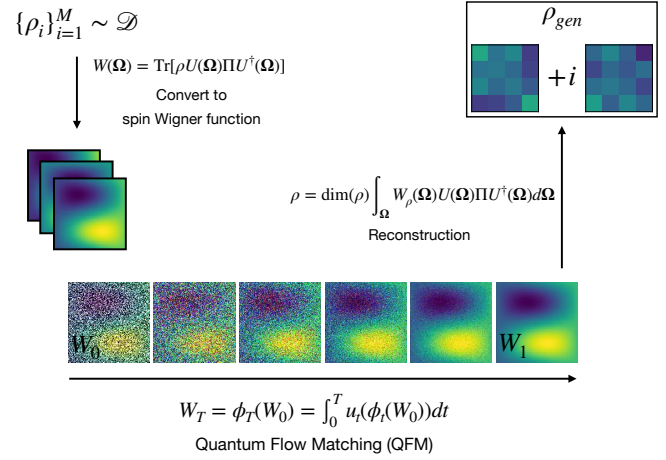


Figure 1. **Overall procedure of QFM.** We bypass the direct learning of quantum states by converting them into spin Wigner function and learn the underlying distribution. This approach allows us to effectively generate physically valid and accurate quantum states, which was not achievable by direct learning from density matrices.

classical data such as images or classical PDEs. These are inadequate for applications requiring consistency with important physical quantities such as purity, entanglement entropy, and quantum phases of matter. Second, complex-valued density matrices pose a significant challenge due to the sign structure problem, making previous methods unsuitable for handling them (Westerhout et al., 2020; Dugan et al., 2023).

To address this challenge, we propose Quantum Flow Matching (QFM), a novel generative modeling method that leverages flow matching to effectively learn the distributions of quantum states. The procedure of QFM, summarized in Figure 1, can be illustrated in the following steps. First, we convert quantum states into informationally complete functional representations using spin Wigner functions. Second, we employ Functional Flow Matching (FFM) to learn the distribution of spin Wigner functions in function space. This allows us to generate new spin Wigner functions that accurately reflect the underlying quantum distribution. Finally, the new quantum states are reconstructed from these generated spin Wigner functions.

Our contributions are summarized as follows:

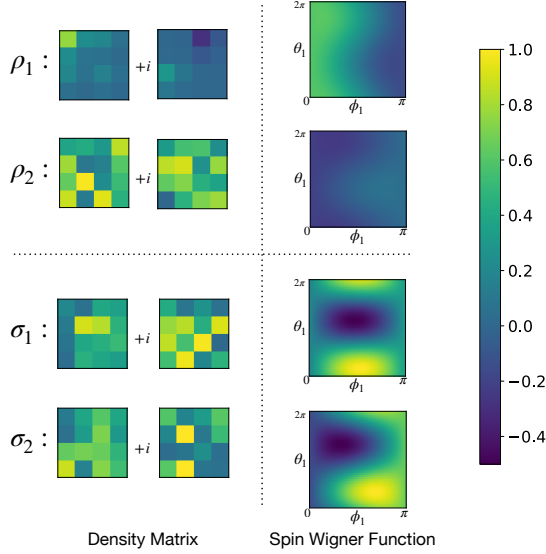


Figure 2. An Example of two-qubit product states ρ_1, ρ_2 (top) and two-qubit maximally entangled states σ_1, σ_2 (bottom), in both density matrix (left) and spin Wigner function (right). For the spin Wigner functions, parameterizations on θ_1 and ϕ_1 are shown, with θ_2 and ϕ_2 fixed at π and $\pi/2$, respectively. Random arbitrary single-qubit rotation gates $U(\alpha, \beta, \gamma) = \exp(-i\alpha Z/2) \exp(-i\beta Y/2) \exp(-i\gamma X/2)$ were applied to both qubits of $|00\rangle$ for the product states and to both qubits of $(|00\rangle + |11\rangle)/\sqrt{2}$ for the maximally entangled states. Unlike in the density matrix form, the product states and maximally entangled states are clearly distinguishable in the spin Wigner function representation.

- We propose QFM, a pioneering generative model for quantum distributions that accurately capture the underlying physics of a given quantum system.
- We bypass the direct density matrix learning by using spin Wigner function. FFM is utilized to leverage resolution-invariant nature of functional models, enabling accurate reconstruction of quantum states from the learned spin Wigner functions.
- We demonstrate that our method effectively learns the underlying physics of given quantum systems by evaluating various physical quantities, such as trace, purity, and entanglement entropy.

2. Learning Quantum Distribution with Flow Matching

2.1. Quantum State to Spin Wigner Function

Learning the distributions of quantum states using deep generative models requires effective classical representations of these quantum states. Training the generative model using the state vector or density matrix representation alone is insufficient, as it does not ensure that important physical quantities are accurately captured.

The Wigner function effectively addresses these shortcomings by mapping a quantum density matrix onto a quasi-probability distribution in classical phase space (Rundle et al., 2017; Wigner, 1932). However, since the Wigner function has been primarily explored in continuous variable and infinite-dimensional Hilbert spaces, such as in quantum optics and chemistry, its application in machine learning is still limited in these settings (Dugan et al., 2023).

In this work, we utilize spin Wigner function, which fully describes quantum systems of arbitrary (including finite) dimensions (Tilma et al., 2016; Rundle et al., 2017). This approach enables the application of deep generative models to a range of intriguing tasks involving finite-dimensional Hilbert spaces, such as generating the ground state of a quantum spin chain with specific quantum phases.

For an arbitrary quantum state ρ , the spin Wigner function is given by $W(\Omega) = \text{Tr}[\rho U(\Omega) \Pi U^\dagger(\Omega)]$, where $U(\Omega)$ and Π are analogous to the displacement and parity operator in the original Wigner function, respectively. With an appropriate choice of Hermitian observable Π , along with unitary U and parameterization Ω , the original density matrix can be fully reconstructed as $\rho = \text{dim}(\rho) \int_{\Omega} W(\Omega) U(\Omega) \Pi U^\dagger(\Omega) d\Omega$, where $\text{dim}(\rho)$ is the dimensionality of ρ .

Here, $d\Omega$ should be chosen such that $U(\Omega)$ follows a Haar random, or unitary 2-design, distribution. Specifically, unitary 2-design, denoted by \mathcal{U} , is a distribution of unitary matrices where the expectation value matches that of the Haar distribution up to the second moment. Then, the integral term of the reconstruction is given by

$$\begin{aligned} \int_{\Omega} W(\Omega) \Delta(\Omega) d\Omega &= \mathbb{E}_{\mathcal{U}} \text{Tr}_A[(\rho \otimes I_B) U^{\otimes 2} \Pi^{\otimes 2} U^{\dagger \otimes 2}] \\ &= \text{Tr}_A[(\rho \otimes I_B)(c_{\Pi, I} I + c_{\Pi, F} F)] = c_{\Pi, I} I + c_{\Pi, F} \rho, \end{aligned}$$

where $\Delta(\Omega) \equiv U(\Omega) \Pi U^\dagger(\Omega)$, and expectation over \mathcal{U} is taken with respect to the Haar distribution. I and F denote identity and swap operator on the composite system, respectively, and the constants are calculated as $c_{\Pi, I} = 0$ and $c_{\Pi, F} = 1/\text{dim}(\rho)$ (Mele, 2024). Therefore, we can fully reconstruct the original density matrix ρ .

To generalize this to N qubits, there is flexibility in the choice of $U(\Omega)$ and Π . Throughout our work, we adhere to the conventions of Rundle et al. (2017), where Π_N is $2^N \times 2^N$ diagonal matrix with $2^{-N}[1 + (2^N - 1)\sqrt{2^N + 1}]$ as the first element and $2^{-N}[1 - \sqrt{2^N + 1}]$ as the remaining elements. This choice aligns with the previous calculations of $c_{\Pi, I}$ and $c_{\Pi, F}$. The unitary operator and parameterization are generalized to $U_N(\Omega) = \bigotimes_{j=1}^N e^{i\theta_j Z_j} e^{i\phi_j Y_j}$, where $\Omega \equiv (\vec{\theta}, \vec{\phi})$ and $Z_j(Y_j)$ denotes the Pauli $Z(Y)$ operator on j -th qubit.

Specifically, we reconstruct N -qubit density matrix ρ from

spin Wigner Function by the following approximation:

$$\begin{aligned}\rho &= \dim(\rho) \mathbb{E}_{\mathcal{U}} [W(\Omega) \Delta(\Omega)] \\ &= \dim(\rho) \int_{\otimes^N S^2} W(\Omega) \Delta(\Omega) \prod_{i=1}^N \sin \phi_i d\theta_1 d\phi_1 \cdots d\theta_N d\phi_N \\ &\simeq \pi^N \cdot \mathbb{E}_{\substack{\theta_i \sim U[0, 2\pi] \\ \phi_i \sim U[0, \pi]}} \left[W(\Omega) \Delta(\Omega) \prod_{i=1}^N \sin \phi_i \right].\end{aligned}$$

Using this approach, we can construct an informationally complete representation of N -qubit density matrices with functions of $2N$ parameters. Figure 2 demonstrates the two-qubit product states and maximally entangled states in the spin Wigner representation.

2.2. Functional Flow Matching

Flow Matching (FM) is a continuous normalizing flow that models the integration of vector field u_t , or the flow ϕ_t , by learning the vector field of the prescribed dynamics. This can be described with the ordinary differential equation $\frac{d}{dt} \phi_t(W_t) = u_t(\phi_t(W_t))$, where W_t is the spin Wigner function at time $t \in [0, 1]$. Lipman et al. (2022) proposed conditional flow matching objective (CFM), that learns the marginal vector field $u_t(W_t)$ which corresponds to the target probability density path $p_t(W_t)$, known to be more stable and robust compared to diffusion models (Ho et al., 2020; Song et al., 2020). The CFM objective is given by

$$\mathcal{L}_{\text{CFM}} = \mathbb{E}_{t \sim U(0,1), W_t \sim p_t(W_t)} \|v_t(W_t) - u_t(W_t|W_1)\|^2,$$

where $v_t(W)$ is a trainable neural network.

To learn the distribution of spin Wigner function, we utilize Functional Flow Matching, which is a functional version of FM (Kerrigan et al., 2023), generalized to the function space, implemented with Fourier Neural Operator (Li et al., 2020). It features a resolution-invariant property of functional models, and we leverage it when reconstructing the generated spin Wigner functions back into the quantum states.

3. Experiments

Single Qubit: Trace, Purity. We first demonstrate the generation of single-qubit density matrices using QFM. Our objective is to generate quantum states with a specified purity $\alpha = \text{Tr}[\rho^2]$, where ρ denotes density matrix of the quantum state. Moreover, we also examine the trace of the generated data to ensure it satisfies the unit trace constraint.

We conduct experiments with a synthetic dataset consisting of 7000 pure states ($\alpha = 1$) and mixed states ($\alpha \in \{0.625, 0.905\}$). Single-qubit quantum states can be expressed as $(I + \mathbf{n} \cdot \boldsymbol{\sigma})/2$, where \mathbf{n} and $\boldsymbol{\sigma}$ are Bloch and Pauli vectors, respectively. Datasets of purity α are prepared by

Table 1. Quantitative comparisons of QFM and direct learning of FM on a single qubit. α , $\text{Tr}[\rho]$ and $\text{Tr}[\rho^2]$ represent the purity of training dataset, average values of the trace, and the purity over the generated quantum states, respectively.

Method	α	$\text{Tr}[\rho]$	$\text{Tr}[\rho^2]$
FM	1	0.998	0.714
QFM	1	0.998	1.004
FM	0.625	1.001	0.531
QFM	0.625	0.997	0.625
FM	0.905	0.999	0.622
QFM	0.905	0.995	0.895

Table 2. Quantitative comparisons of QFM and direct learning of FM on multi qubits. S , $\text{Tr}[\rho]$ and $S(\rho_A)$ represent the entanglement entropy of training dataset, the average values of the trace, and the entanglement entropy over the generated quantum states, respectively.

Method	S	$\text{Tr}[\rho]$	$S(\rho_A)$
FM	0.7	1.009	0.783
QFM	0.7	0.934	0.694

sampling \mathbf{n} from uniform distribution with the constraint $\|\mathbf{n}\|_2 = \sqrt{2\alpha - 1}$.

Multi qubits: Trace, Entanglement Entropy. We further demonstrate the capability of QFM in a multi-qubit setting. Specifically, we focus on generating 2-qubit quantum states with a desired entanglement entropy $S = -\text{Tr}[\rho_A \log \rho_A]$, where ρ_A is the reduced density matrix. We prepare a synthetic dataset consisting of 900 quantum states with fixed entanglement entropy $S = 0.7$. This is done by initializing two-qubit quantum states with an entanglement entropy of 0.7 and applying random arbitrary single-qubit rotations.

Tables 1 and 2 present a comparative analysis between QFM and direct learning of FM using density matrices. QFM effectively captures the physical quantities of the dataset compared to direct learning of FM. This is demonstrated by its better performance in the target purity values (Table 1) and entanglement entropy (Table 2).

4. Future Work

One potential future direction is extending the method to larger multi-qubits system such as generating genuinely multipartite entangled states (Goyeneche & Życzkowski, 2014) and the ground states of generalized cluster Hamiltonian with specific quantum phases (Caro et al., 2022; Gil-Fuster et al., 2024). Another interesting direction is quantum error mitigation by leveraging generative model as denoiser and formulating as solving an inverse problem (Kawar et al., 2022; Chung et al., 2022; Shim & Kim, 2024). Lastly, by simultaneously utilizing Quantum Autoencoders (Romero et al., 2017), our method potentially enables the accurate generation of high-dimensional quantum states, which have been intractable due to heavy computational costs.

References

- Carleo, G., Cirac, I., Cranmer, K., Daudet, L., Schuld, M., Tishby, N., Vogt-Maranto, L., and Zdeborová, L. Machine learning and the physical sciences. *Reviews of Modern Physics*, 91(4):045002, 2019.
- Caro, M. C., Huang, H.-Y., Cerezo, M., Sharma, K., Sornborger, A., Cincio, L., and Coles, P. J. Generalization in quantum machine learning from few training data. *Nature communications*, 13(1):4919, 2022.
- Carrasquilla, J., Torlai, G., Melko, R. G., and Aolita, L. Reconstructing quantum states with generative models. *Nature Machine Intelligence*, 1(3):155–161, 2019.
- Chung, H., Sim, B., Ryu, D., and Ye, J. C. Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*, 35:25683–25696, 2022.
- Dugan, O. M., Lu, P. Y., Dangovski, R., Luo, D., and Soljagic, M. Q-flow: generative modeling for differential equations of open quantum dynamics with normalizing flows. In *Proc. of the International Conference on Machine Learning (ICML)*, 2023.
- Gil-Fuster, E., Eisert, J., and Bravo-Prieto, C. Understanding quantum machine learning also requires rethinking generalization. *Nature Communications*, 15(1):2277, 2024.
- Goyeneche, D. and Życzkowski, K. Genuinely multipartite entangled states and orthogonal arrays. *Physical review A*, 90(2):022316, 2014.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. In *Advances in neural information processing systems (NeurIPS)*, 2020.
- Kawar, B., Elad, M., Ermon, S., and Song, J. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022.
- Kerrigan, G., Migliorini, G., and Smyth, P. Functional flow matching. *arXiv:2305.17209*, 2023.
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A. Fourier neural operator for parametric partial differential equations. *arXiv:2010.08895*, 2020.
- Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling. *arXiv:2210.02747*, 2022.
- Mele, A. A. Introduction to haar measure tools in quantum information: A beginner’s tutorial. *Quantum*, 8:1340, 2024.
- Romero, J., Olson, J. P., and Aspuru-Guzik, A. Quantum autoencoders for efficient compression of quantum data. *Quantum Science and Technology*, 2(4):045001, 2017.
- Rundle, R. P., Mills, P., Tilma, T., Samson, J., and Everitt, M. J. Simple procedure for phase-space measurement and entanglement validation. *Physical Review A*, 96(2):022117, 2017.
- Shim, J. Y. and Kim, J. Diffusion-based quantum error mitigation using stochastic differential equation. *arXiv preprint arXiv:2405.14283*, 2024.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. *arXiv:2011.13456*, 2020.
- Tilma, T., Everitt, M. J., Samson, J. H., Munro, W. J., and Nemoto, K. Wigner functions for arbitrary quantum systems. *Physical review letters*, 117(18):180401, 2016.
- Westerhout, T., Astrakhantsev, N., Tikhonov, K. S., Katsnelson, M. I., and Bagrov, A. A. Generalization properties of neural network approximations to frustrated magnet ground states. *Nature communications*, 11(1):1593, 2020.
- Wigner, E. On the quantum correction for thermodynamic equilibrium. *Physical review*, 40(5):749, 1932.