

Exploring the flavor structure of leptons via diffusion models

Satsuki Nishimura^{1,*}, Hajime Otsuka^{1,2,†} and Haruki Uchiyama^{1,‡}

¹*Department of Physics, Kyushu University, 744 Motoooka, Nishi-ku, Fukuoka 819-0395, Japan*

²*Quantum and Spacetime Research Institute (QuaSR), Kyushu University,
744 Motoooka, Nishi-ku, Fukuoka 819-0395, Japan*



(Received 11 October 2025; accepted 23 February 2026; published 18 March 2026)

We propose a method to explore the flavor structure of leptons using diffusion models, which are known as one of generative artificial intelligence (generative AI). We consider a simple extension of the Standard Model with the type I seesaw mechanism and train a neural network to generate the neutrino mass matrix. By utilizing transfer learning, the diffusion model generates 104 solutions that are consistent with the neutrino mass squared differences and the leptonic mixing angles. The distributions of the CP phases and the sums of neutrino masses, which are not included in the conditional labels but are calculated from the solutions, exhibit nontrivial tendencies. In addition, the effective mass in neutrinoless double beta decay is concentrated near the boundaries of the existing confidence intervals, allowing us to verify the obtained solutions through future experiments. An inverse approach using the diffusion model is expected to facilitate the experimental verification of flavor models from a perspective distinct from conventional analytical methods.

DOI: [10.1103/rtn-d-vwt9](https://doi.org/10.1103/rtn-d-vwt9)

I. INTRODUCTION

The flavor structure of quarks and leptons is one of the intriguing puzzles in the Standard Model of particle physics. Current experimental and observational data suggest the existence of massive neutrinos and large mixing angles in the lepton sector.

The flavor structure of neutrinos has been addressed in previous works using both top-down and bottom-up approaches. In the top-down approach, we assume specific textures for the Yukawa couplings of charged leptons and neutrinos, which originate from continuous symmetries, such as a $U(1)$ flavor symmetric model using the Froggatt-Nielsen (FN) mechanism [1], non-Abelian discrete symmetries (see for reviews, Refs. [2–8]), modular flavor symmetries [9] (see for reviews, Refs. [10,11]) and selection rules without group actions [12,13]. These approaches lead to various textures of neutrino mass matrices at the low-energy scale. Based on the derived textures of Yukawa couplings, one can exhibit the phenomenologically viable distributions for observables such as active neutrino

masses, size of CP violation, and lepton flavor violation. This top-down approach is useful to discuss the connection between low-energy observables and the predictions in flavor symmetric models. On the other hand, in the bottom-up approach as adopted in, for example, Refs. [14–16], they find out a form of neutrino Yukawa couplings that reproduces the experimental data at the low-energy scale. Then, they discuss the lepton flavor violation such as $\text{Br}(\mu \rightarrow e\gamma)$ in the framework of supersymmetric models. In addition, Ref. [17] discusses baryon asymmetry based on existing experimental results by inverting formula of the seesaw mechanism.

The recent developments of machine learning have been utilized in the top-down approach to search for physics beyond the Standard Model, e.g., the flavor structure of quarks and leptons in the $U(1)$ Froggatt-Nielsen model with reinforcement learning [18], the phenomenology of axion model from flavor [19]. In this paper, we adopt a bottom-up approach utilizing generative artificial intelligence (generative AI) to explore the flavor structure of leptons. In particular, we focus on a diffusion model, which is one of the generative models that has been used in high-energy physics, for example, in simulating electron-proton scattering events at the future Electron-Ion Collider [20]. In the context of denoising diffusion probabilistic models (DDPMs) [21], random Gaussian noise is progressively added to a given training data until it approaches pure noise. Subsequently, a neural network is constructed to predict the actual noise that has been added, using these noisy data as input. After the training process, the diffusion

* Contact author: nishimura.satsuki@phys.kyushu-u.ac.jp

† Contact author: otsuka.hajime@phys.kyushu-u.ac.jp

‡ Contact author: uc.haruki496ym@gmail.com

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

model generates new data by denoising through the trained neural network. Furthermore, it is possible to generate desired data from the initial dataset by applying arbitrary conditional labels. This class of diffusion models is referred to as conditional diffusion models.

We employ the DDPM with classifier-free guidance, specifically one of the conditional diffusion models, to investigate the unknown flavor structure of neutrinos within a basis in which the Yukawa matrix of charged leptons and gauge interactions are flavor diagonal. In particular, we study the Standard Model with three right-handed neutrinos where the light neutrino masses are generated through the seesaw mechanism [22–25]. Our training dataset consists of randomly chosen neutrino Yukawa couplings and right-handed neutrino masses, similar to the analysis of neutrino mass anarchy [26]. The labels are defined as the neutrino mass squared differences and mixing angles in the lepton sector. After learning the diagonalization process of the light neutrino mass matrix, we fix the labels to the observed neutrino mass squared differences and mixing angles in the lepton sector at the 3σ level. Consequently, we can obtain a sufficient number of neutrino Yukawa couplings and right-handed neutrino masses that reproduce the observed data. These results lead to reveal the phenomenologically viable distributions of observables such as the effective mass in neutrinoless double beta decay and leptonic CP violation. Although an inverse approach for reconstructing Yukawa matrices from low-energy neutrino data has been discussed in the literature (see, e.g., Refs. [16,17]), we present a new method that employs the diffusion model as a sampler.

This paper is organized as follows. After presenting our setup in Sec. II A, the conditional diffusion model is introduced in Sec. II B. In this section, we explain the diffusion process to train the neural network in Sec. II B 1 and the reverse process to generate new data in Sec. II B 2. In addition, we perform transfer learning in Sec. II B 3 to improve the consistency of the generated samples with experimental data. We show results generated by the diffusion model in Sec. III. Sec. IV is devoted to the conclusion and future prospects. In Appendix A, we give the formulation of diffusion models.

Furthermore, we explain the details of conditional diffusion models in Appendix B and transfer learning in Appendix C.

II. FLAVOR STRUCTURE OF LEPTONS VIA THE CONDITIONAL DIFFUSION MODELS

After summarizing the flavor structure of leptons in Sec. II A, we implement the neutrino mass matrices to the conditional diffusion models in Sec. II B.

A. Preliminaries

We focus on the Majorana neutrino mass with the seesaw mechanism. The relevant Lagrangian is

$$\mathcal{L} = Y_{i\alpha}^\nu \bar{N}_i \ell_\alpha H - \frac{1}{2} M_{ij} \bar{N}_i \bar{N}_j + \text{H.c.}, \quad (2.1)$$

where ℓ_α , N_j , and H respectively denote the left-handed lepton doublet, the Higgs doublet, and the right-handed neutrinos. $Y_{i\alpha}^\nu$ and M_{ij} respectively represent Yukawa matrix for neutrinos and Majorana mass matrix for right-handed neutrinos in a basis where the charged lepton Yukawa matrix is diagonal. α denote the lepton flavor indices $\alpha = e, \mu, \tau$. By integrating out heavy $\bar{N}_i$, the mass matrix for active neutrinos is given by

$$(m_\nu)_{\alpha\beta} = \langle H \rangle^2 (Y^\nu)_{\alpha i} (M^{-1})_{ij} (Y^\nu)_{j\beta}, \quad (2.2)$$

with $\langle H \rangle$ being the vacuum expectation value of the Higgs field. A complex-valued symmetric matrix m_ν can be diagonalized by a unitary matrix U_{PMNS}

$$U_{\text{PMNS}}^T m_\nu U_{\text{PMNS}} = \text{diag}(m_1, m_2, m_3), \quad (2.3)$$

where $\{m_1, m_2, m_3\}$ are taken as real and positive values. The neutrino mass ordering is given by $m_1 < m_2 < m_3$ for the normal hierarchy and $m_3 < m_1 < m_2$ for the inverted hierarchy.

The mixing matrix U_{PMNS} can be parametrized by three mixing angles θ_{ij} , a CP -violating Dirac phase δ_{CP} and two Majorana phases α_{21}, α_{31}

$$U_{\text{PMNS}} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\frac{\alpha_{21}}{2}} & 0 \\ 0 & 0 & e^{i\frac{\alpha_{31}}{2}} \end{pmatrix}. \quad (2.4)$$

Here, the mixings $c_{ij} = \cos \theta_{ij}$ and $s_{ij} = \sin \theta_{ij}$ with $i, j = 1, 2, 3$ and $i < j$ are rewritten as

$$\begin{aligned} \sin^2 \theta_{13} &= |(U_{\text{PMNS}})_{13}|^2, & \sin^2 \theta_{23} &= \frac{|(U_{\text{PMNS}})_{23}|^2}{1 - |(U_{\text{PMNS}})_{13}|^2}, \\ \sin^2 \theta_{12} &= \frac{|(U_{\text{PMNS}})_{12}|^2}{1 - |(U_{\text{PMNS}})_{13}|^2}. \end{aligned} \quad (2.5)$$

A rephasing-invariant measure of CP violation is given by the Jarlskog invariant, J_{CP} from PMNS matrix elements

$$J_{CP} = \text{Im}[U_{e1}U_{\mu 2}U_{e2}^*U_{\mu 1}^*] = s_{23}c_{23}s_{12}c_{12}s_{13}c_{13}^2 \sin \delta_{CP}, \quad (2.6)$$

with $U_{ai} := (U_{\text{PMNS}})_{ai}$, and the Majorana phases are also estimated by other invariants I_1 and I_2

TABLE I. Experimental values for the neutrino mass differences, mixing angles and CP phase obtained from global analysis of NuFIT 6.0 with Super-Kamiokande atmospheric data [29], where $\Delta m_{3l}^2 \equiv \Delta m_{31}^2 = m_3^2 - m_l^2 > 0$ for NO and $\Delta m_{3l}^2 \equiv \Delta m_{32}^2 = m_3^2 - m_l^2 < 0$ for IO.

Observables	Normal ordering (NO)		Inverted ordering (IO)	
	1 σ range	3 σ range	1 σ range	3 σ range
$\sin^2 \theta_{12}$	$0.308^{+0.012}_{-0.011}$	$0.275 \rightarrow 0.345$	$0.308^{+0.012}_{-0.011}$	$0.275 \rightarrow 0.345$
$\sin^2 \theta_{13}$	$0.02215^{+0.00056}_{-0.00058}$	$0.02030 \rightarrow 0.02388$	$0.02231^{+0.00056}_{-0.00056}$	$0.02060 \rightarrow 0.02409$
$\sin^2 \theta_{23}$	$0.470^{+0.017}_{-0.013}$	$0.435 \rightarrow 0.585$	$0.550^{+0.012}_{-0.015}$	$0.440 \rightarrow 0.584$
δ_{CP}/π	$1.18^{+0.14}_{-0.23}$	$0.69 \rightarrow 2.02$	$1.52^{+0.12}_{-0.14}$	$1.12 \rightarrow 1.86$
$\frac{\Delta m_{21}^2}{10^{-5} \text{ eV}^2}$	$7.41^{+0.21}_{-0.20}$	$6.82 \rightarrow 8.03$	$7.42^{+0.21}_{-0.20}$	$6.82 \rightarrow 8.04$
$\frac{\Delta m_{3l}^2}{10^{-3} \text{ eV}^2}$	$2.507^{+0.026}_{-0.027}$	$2.427 \rightarrow 2.590$	$-2.486^{+0.025}_{-0.028}$	$-2.570 \rightarrow -2.406$

$$I_1 = \text{Im}[U_{e1}^* U_{e2}] = c_{12} s_{12} c_{13}^2 \sin\left(\frac{\alpha_{21}}{2}\right),$$

$$I_2 = \text{Im}[U_{e1}^* U_{e3}] = c_{12} s_{13} c_{13} \sin\left(\frac{\alpha_{31}}{2} - \delta_{CP}\right). \quad (2.7)$$

Experimental values for the neutrino mass squared differences, mixing angles, and Dirac CP violation are summarized in Table I and the PMNS matrix at the 3 σ CL range is of the form

$$|U_{\text{PMNS,exp}}|_{3\sigma} = \begin{pmatrix} 0.801 \rightarrow 0.842 & 0.519 \rightarrow 0.580 & 0.142 \rightarrow 0.155 \\ 0.252 \rightarrow 0.501 & 0.496 \rightarrow 0.680 & 0.652 \rightarrow 0.756 \\ 0.276 \rightarrow 0.518 & 0.485 \rightarrow 0.673 & 0.637 \rightarrow 0.743 \end{pmatrix}. \quad (2.8)$$

The sum of neutrino masses is constrained as $\sum_i m_i < 0.12 \text{ eV}$ (95% CL) from the data of cosmic microwave background (CMB) [27] which will be further constrained as $0.059 \text{ eV} < \sum_i m_i < 0.113 \text{ eV}$ at 95% CL [28] by combining the data from the Dark Energy Spectroscopic Instrument with CMB data.

The effective Majorana neutrino mass $m_{\beta\beta}$ for the neutrinoless double beta decay is given by

$$\langle m_{\beta\beta} \rangle = |m_1 c_{12}^2 c_{13}^2 + m_2 s_{12}^2 c_{13}^2 e^{i\alpha_{21}} + m_3 s_{13}^2 e^{i(\alpha_{31} - 2\delta_{CP})}|, \quad (2.9)$$

which is upper bounded by 0.036 eV (90% CL) [30]. Moreover, the effective mass of electron neutrino probed by tritium beta decay is defined as follows:

$$m_{\nu_e} = (m_1^2 c_{12}^2 c_{13}^2 + m_2^2 s_{12}^2 c_{13}^2 + m_3^2 s_{13}^2)^{1/2}, \quad (2.10)$$

which is upper bounded by 0.45 eV (90% CL) [31].

B. Conditional diffusion models

In this section, we introduce the conditional diffusion model used in our analysis. That is constructed of a basic diffusion model and classifier-free guidance (CFG). We refer the reader to Appendix A for a formulation of the diffusion models based on denoising diffusion probabilistic models (DDPMs) and to Appendix B for the detail of the CFG.¹

The diffusion model consists of two stages, referred to as the diffusion process and the reverse process. In the diffusion process, random noise ϵ is added to input data G , and a machine learns to predict the added noise. In the reverse process, noise ϵ_θ output by the machine is gradually removed from the pure Gaussian noise to produce meaningful data. To predict the noise, we utilize neural network models. In a neural network, a n th layer with N_{n-1} dimensional vector $\vec{X}_{n-1} = (X_{n-1}^1, X_{n-1}^2, \dots, X_{n-1}^{N_{n-1}})$ transforms into a N_n dimensional vector $\vec{X}_n = (X_n^1, X_n^2, \dots, X_n^{N_n})$ through the following way:

$$X_n^i = h_n(w_n^{ij} X_{n-1}^j + b_n^i). \quad (2.11)$$

Here, h is the activation function, w is the weight, and b is the bias, respectively. Since the activation function is generally a nonlinear function, the neural network acts on multiple nonlinear transformations. In our analysis, fully

¹This study demonstrates the effectiveness of generative AI, particularly diffusion models, in analyzing concrete flavor models. While DDPM was adopted here as one approach providing the fundamental formulation for diffusion models, there are other methods alongside technological advancements, such as conditional flow matching (CFM) [32], ODE Solvers for diffusion models [33], and hybrid DDPM-CFM architectures. In addition, Ref. [34] provides a review of various techniques on diffusion models. Achieving fast and stable data generation using these state-of-the-art methods remains a future challenge for enabling more efficient analysis.

connected layers are considered. In the following, we implement the Standard Model with three right-handed neutrinos to the diffusion model, where the diffusion and reverse processes are respectively described in Secs. II B 1 and II B 2.

1. Diffusion process

In the diffusion process, let us define the data sequence $\{x_1, x_2, \dots, x_T\}$ associated with an initial data x_0

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{\bar{\beta}_t}\epsilon, \quad (2.12)$$

with $t = 1, \dots, T$. Here, ϵ obeying a standard normal distribution $\mathcal{N}(0, 1)$ is called a noise, and $\bar{\alpha}_t, \bar{\beta}_t$ are determined as follows:

$$\bar{\alpha}_t = \prod_{s=1}^t \alpha_s, \quad \bar{\beta}_t = 1 - \bar{\alpha}_t, \quad (2.13)$$

with

$$\alpha_t = 1 - \beta_t, \quad \beta_t = \left(1 - \frac{t}{T}\right)\beta_{\min} + \frac{t}{T}\beta_{\max}, \quad (2.14)$$

where α_t, β_t are called noise schedules.² Throughout our analysis, we adopt $\beta_{\min} = 10^{-4}$, $\beta_{\max} = 0.02$, and $T = 1000$. We also tested cosine schedule [35] and exponential schedule for β_t , but ultimately adopted the linear schedule to ensure generation stability.

We consider the diagonal basis of Majorana mass matrix such as $M = \text{diag}(M_1, M_2, M_3)$, where M_1, M_2 are complex numbers and M_3 is a real number. Then, in preparing the initial data, we deal with the following values:

$$G = \left\{ \text{Re } Y_{i\alpha}^\nu, \text{Im } Y_{i\alpha}^\nu, \frac{|M_1|}{M_3}, \frac{|M_2|}{M_3}, \log_{10} \left(\frac{\langle H \rangle^2}{M_3} [\text{meV}] \right), \arg M_1, \arg M_2 \right\}, \quad (2.15)$$

$$L = \{ \log_{10}(\Delta m_{21}^2 [\text{meV}^2]), \log_{10}(\Delta m_{31}^2 [\text{meV}^2]), |(U_{\text{PMNS}})_{ij}| \}, \quad (2.16)$$

with $i, j = 1, 2, 3$. G is a main input in the diffusion model and indicates the set of variables that are subject to data generation. On the other hand, L is an auxiliary label for conditional generation and is passed to the neural network in the appropriate way described later. Now, G has 23 components and L has 11 components in terms of real numbers. Each element of G was generated by uniform random numbers within the following ranges:

$$-1 \leq \{ \text{Re } Y_{i\alpha}^\nu, \text{Im } Y_{i\alpha}^\nu \} \leq 1, \quad 10^{-5} \leq \left\{ \frac{|M_1|}{M_3}, \frac{|M_2|}{M_3} \right\} \leq 1, \quad (2.17)$$

$$-\pi \leq \{ \arg M_1, \arg M_2 \} \leq \pi, \quad -1.22 \leq \log_{10}(\langle H \rangle^2 / M_3 [\text{meV}]) \leq 2.78. \quad (2.18)$$

In particular, Eq. (2.18) means $6.05 \times 10^{-2} \text{ meV} \leq \langle H \rangle^2 / M_3 \leq 6.05 \times 10^2 \text{ meV}$, i.e., $10^{14} \text{ GeV} \leq M_3 \leq 10^{18} \text{ GeV}$. Moreover, we imposed $|M_1|/M_3 \leq |M_2|/M_3$. Then, L is calculated based on Eq. (2.3) from a random generated G .³ In actual learning, we prepared 100,000 pairs of (G, L) . Of these data, 90% are used as training data and 10% as validation data.

The training flow of the neural network in our diffusion model is shown in Fig. 1. When the integer t is randomly selected from $[1, T]$, the noisy data x_t is determined according to Eq. (2.12). When an input layer of the neural network receives noisy pair (x_t, L) as input X_0 in Eq. (2.11), it performs nonlinear transformations according to that equation. An output of the network ϵ_θ is calculated as the predicted noise, and an error with the added noise ϵ is evaluated by a loss function. The parameters $\theta = w, b$ of the network are updated by back-propagation based on the value of the loss function, and a well-trained network can accurately estimate the noise component in x_t .

²Among the various types of generative AI, the diffusion models with CFG are superior in that it can easily and accurately generate data based on conditional labels L . In the context of image generation, various paintings are collected as input data G and associated information as L is also prepared. The inputs G correspond to added noises according to the noise schedules. Based on these data, a machine that has learned to predict noise will be able to output motifs and painting styles specific to an artist. As a result, for example, given a command “make an image of ‘The Persistence of Memory’ by Salvador Dalí with a touch of Van Gogh”, the title of art or the names of artists function as conditional labels L , and an image that combines the motif “melting clock” and Van Gogh’s “The starry night” is output.

³The CP phase δ_{CP} is not included in the conditional label L due to its large experimental uncertainty. More generally, CP -sensitive quantities such as J_{CP} and Majorana invariants were not included among the conditioning variables. This omission is a deliberate modeling choice, rather than a technical limitation. By restricting the conditioning variables to the absolute values of PMNS matrix, we aimed to keep the conditioning minimal and to avoid implicitly injecting CP -phase information into the generative process. This setup allows the CP phases to be treated as outcomes of the model and enables an exploratory examination of their distributions under the requirement of reproducing experimentally established observables.




FIG. 1. The summary of input/output of a neural network in the diffusion process. Based on the noise schedule, noise ϵ sampled from the standard normal distribution $\mathcal{N}(0, 1)$ is added to the data G . The inputs of the network are the noisy data x_t and the label information $\{L, t\}$, while the output is the predicted noise ϵ_θ . In particular, $\{L, t\}$ is also input to the intermediate layers for conditional learning. The network is updated to minimize a mean squared error (MSE) defined from the actual added noise ϵ and predictive noise ϵ_θ .

To make a conditional diffusion model, we give hidden layers the labels L and the time t in probability of 90%. On the other hand, in 10%, we dropped out the labels by setting $L = \emptyset$. In our program, zeros are input to the neural network as $\emptyset$.

The detailed architecture of our neural network is shown in Table II. The activation function h in Eq. (2.11) is chosen as a SELU function for hidden layers H1 to H6, a tanh function for a hidden layer H7, and an identity function for an output layer. Since the diffusion model is used to predict noise sampled from a normal distribution, H7 contributes to reduce the divergence by limiting the range of output values with tanh. In addition, a loss function is chosen as the MSE, and we use the ADAM optimizer in PyTorch to update the weights w and biases b in Eq. (2.11). A learning rate is automatically adjusted through the scheduler “OneCycleLR” provided by PyTorch, and we adopt $r_{\max} = 0.001$ as a maximum rate. Then, a batch size is 64 and an updating of the parameter is repeated 100,000 times.

2. Reverse process

In the reverse process, the generation of the dataset is performed with labels that accurately reflect reality. Specifically, L is designated based on the central values in Eq. (2.8) and Table I

$$\log_{10}(\Delta m_{21}^2 [\text{meV}^2]) = \log_{10}(7.49 \times 10^1) = 1.87, \quad (2.19)$$

$$\log_{10}(\Delta m_{31}^2 [\text{meV}^2]) = \log_{10}(2.51 \times 10^3) = 3.40, \quad (2.20)$$

$$|U_{\text{PMNS}}| = \begin{pmatrix} 0.822 & 0.550 & 0.149 \\ 0.377 & 0.588 & 0.704 \\ 0.397 & 0.579 & 0.690 \end{pmatrix}. \quad (2.21)$$

Here, we consider the structure of neutrino masses as the normal ordering. The diffusion model can generate data based on the inverted ordering by adjusting L to the experimental values for that scenario, but we only focus on the normal ordering in this paper for simplicity. We will hope to report on this analysis for future work.

From the generated G , the label L is recalculated, and the results can be compared to the experimental values. If the obtained result is within a certain error range as described later, the data G reproducing the observables has been generated solely from the experimental results. In this study, the Yukawa matrices and Majorana masses in Eq. (2.15) are generated by referencing experimental results at the electroweak scale. When necessary, these values at a high energy scale can be obtained by appropriately transforming the results obtained by the reverse process according to the renormalization group equations. Note that renormalization group effects may quantitatively influence mixing angles and CP phases, so the distribution presented in the next section should be interpreted with this limitation in mind.

3. Transfer learning

When the diffusion process has been completed once using entirely random training data, various new data G can be generated through the reverse process. However, since

TABLE II. In the neural network, the inputs are the noised data x_t (23 dimensions), the label L (11 dimensions), and the time t . The activation functions are chosen as the SELU function for hidden layers H1 to H6. On the other hand, H7 has the tanh function and the output layer has the identity function. Note that the label L is treated as $\emptyset$ with 10% probability.

Component	Layers	Data dim.	Conditioning dim.	Activation
Input	...	$\mathbb{R}^{23}$		...
Scale up	H1, H2, H3	$\mathbb{R}^{64}, \mathbb{R}^{128}, \mathbb{R}^{256}$		SELU
Deepest hidden	H4	$\mathbb{R}^{512}$	$\mathbb{R}^{12}$	SELU
Scale down	H5, H6	$\mathbb{R}^{256}, \mathbb{R}^{128}$		SELU
Transition	H7	$\mathbb{R}^{64}$		tanh
Output	...	$\mathbb{R}^{23}$	...	Identity

the original random data are inadequate for accurately reproducing the experimental observables, the current network cannot achieve sufficient accuracy. To improve the accuracy of the reproduction of experimental values by the diffusion model, we perform the transfer learning. A detailed explanation of the transfer learning is provided in Appendix C.

After the data G is generated by the diffusion model, the physical values P_ℓ calculated by G can be obtained. The accuracy of P_ℓ , in comparison to the target label L constructed by experimental values, is quantitatively evaluated by the χ^2 function defined as follows:

$$\chi^2 = \sum_{\ell=1}^n \left(\frac{P_\ell - \mu_\ell}{\sigma_\ell} \right)^2, \quad (2.22)$$

where P_ℓ denotes the model-predicted value of a physical observable, μ_ℓ is the central value and σ_ℓ is the corresponding 1σ deviation.⁴

Now, we refer to the first neural network which has been trained once as a *prenetwork*. For the transfer learning, we collect the data that satisfy $\chi^2 < 5.5 \times 10^3$ using the prenetwork. Note that this condition alone would include a lot of data in which both the mass squared differences of neutrinos and mixing angles have low accuracy. In order to improve the efficiency of learning, we collect data in which at least one of the two has high accuracy. Therefore, the following two conditions are taken into account:

- (1) Mass condition: $\chi^2 < 10^3$ is satisfied only for $P_\ell = \{\Delta m_{21}^2, \Delta m_{31}^2\}$.
- (2) Mixing condition: $\chi^2 < 4.5 \times 10^3$ is satisfied only for $P_\ell = \{|(U_{\text{PMNS}})_{ij}|\}$.

We respectively prepare 7,289 models and 36,529 models that satisfy the mass condition and the mixing condition. The number of models satisfying both conditions is counted as 78,488. Thus, a new training data is constructed from the 116,306 models. Based on this new data, the prenetwork was trained once again. Here, hyperparameters are the same as those of the first learning. Additionally, all parameters of the network are updated, so this second training phase is referred to as fine-tuning. In this work, the second network constructed from the prenetwork is referred to as a *tuned-network*.

III. RESULTS

Let us use the tuned-network to generate 3×10^7 data G . Then, we compute the corresponding label L for each of

these generated results and compare it to the experimental values in Eq. (2.8) and Table I. It turns out that there exist 104 solutions such that all physical quantities in L are within the 3σ confidence interval. Note that the solutions obtained here are combinations of numbers that were not included in the training data and that the diffusion model does not select plausible candidates from the random data prepared in advance. In other words, the diffusion model is autonomously generating a new G that reproduces the real observables given as the label L .

We check how much hierarchy occurred in $\{\text{Re } Y_{ia}^\nu, \text{Im } Y_{ia}^\nu\}$ in the generated data G . There are 49, 50, 3, and 2 Yukawa matrices with a total of 18 values having 1, 2, 3, and 4 digits of hierarchy, respectively. This means that 95.2% of the data are within a 2-digit hierarchy. Indeed, in the initial training data in which uniform random numbers were generated in the range of $[-1, 1]$, the percentage of data that fell within a hierarchy of 2 digits was about 94.8%. Therefore, it can be considered that even in the tuned-network that have undergone transfer learning, the uniformity of the initial data is reflected in the generating process. Of course, it is difficult to reproduce experimental values for the mass squared differences of neutrinos and mixing angles within 3σ simply by generating random numbers. By utilizing the diffusion model, Yukawa matrices can be generated that selectively reproduce the desired physical observables.

Figure 2 shows the distribution of right-handed neutrino masses M_1, M_2, M_3 . Since we impose the relation $M_1 \leq M_2 \leq M_3$, the region which do not satisfy this relation is drawn as the gray background in the figure. Although the initial training data is prepared from a wide range of masses, such as Eqs. (2.17) and (2.18), the masses generated by the tuned-network ranged from 10^{15} GeV to 10^{16} GeV. Given that the Yukawa couplings obtained by the tuned-network follow a distribution similar to that of the initial data used in the prenetwork, it seems natural that the right-handed neutrino masses should also be distributed over a wide range. However, in the diffusion model constructed in this study, it turns out that the machine autonomously generates the mass scale around 10^{16} GeV as a favorable scale to satisfy the existing 3σ constraints. Furthermore, a strong linear correlation appears in the scatter plots of M_1 and M_3 . Hence, the generated values are not simply chosen as random numbers but must be output based on some relations learned by the neural network.

We will now consider what implications these solutions provide for other physical values. Fig. 3 shows the correlation between the Dirac CP phase, the mixing angle θ_{23} and the sum of neutrino masses. The median value of the sum is 60.3 meV, and the solutions found by the diffusion model are clustered around this median value. Moreover, the Dirac CP phase tends to appear around 106 deg and 228 deg as first and third quartiles of the 104 data respectively, while there are few solutions around 0 deg or 180 deg. This indicates that

⁴In preliminary analyses, multiple tests are conducted on neural network architectures, including variations in the number of hidden layers. The current design is selected as a best candidate capable of generating practical χ^2 values. Furthermore, although our CFG scale was set to $\gamma = 8.0$, the preliminary analyses also verified values of $\gamma = \{0.5, 1.0, 3.0, 5.0\}$. These values are adjusted to reproduce the experimental results while ensuring that the generated outcomes do not exhibit extreme bias. Additional explanations regarding the CFG scale are provided in Appendix B.




FIG. 2. The distribution of absolute values of right-handed neutrino masses. The color bar shows χ^2 values with $P_\ell = \{\Delta m_{21}^2, \Delta m_{31}^2, s_{12}^2, s_{23}^2, s_{13}^2\}$ in Eq. (2.22). In addition, the white region is allowed by the relation $M_1 \leq M_2 \leq M_3$, but the gray region does not satisfy it.




FIG. 3. The distribution of the Dirac CP phase δ_{CP} with respect to sum of neutrino masses in the left panel and the mixing angle θ_{23} in the right panel. The color bars show χ^2 values with $P_\ell = \{\Delta m_{21}^2, \Delta m_{31}^2, s_{12}^2, s_{23}^2, s_{13}^2\}$ in Eq. (2.22). The solutions are concentrated around $\delta_{CP} = 106, 228$ [deg] (the first and third quartiles) and $\Sigma_i m_i = 60.3$ [meV] (the median value). Note that for the mixing angle in the right panel, all of the points are located within 3σ range because we extract the cases that satisfy the experimental constraints from the data generated by the diffusion model.

the lepton sectors exhibit a significant CP violation in order to explain the existing experimental results.

We also check the results for the Majorana phases α_{21} , α_{31} . Although $\alpha_{21} = 0$ [deg] is calculated for all of the solutions generated by the diffusion model, α_{31} is found to be distributed over a wide angle. Its distribution is shown in Fig. 4. In particular, few candidates are found near $\alpha_{31} = 0$ [deg]. This suggests that CP violation is easily realized in terms of α_{31} . We also examine correlations with θ_{13} , θ_{23} , and the sum of neutrino masses for α_{31} , but no significant correlations are found. Note that the neural network constructed in this study does not directly learn CP symmetry breaking because it lacks the labels related to CP violation. Therefore, our results exhibit the tendencies of CP phases observed in the network's responses under the constraint of matching $|(U_{PMNS})_{ij}|$.

Figure 5 shows the distribution of the effective Majorana neutrino mass $m_{\beta\beta}$ as functions of the sum of neutrino

masses and the electron neutrino mass. The region of the 95% CL with an assumption of the normal ordering is shown with a white background based on NuFIT 6.0 (Ref. [29]), and all the solutions generated by the diffusion model fall within this range. In particular, the red lines indicate the boundaries of the 95% CL, and the data are plotted along that outline. Hence, these solutions could be verified through future experiments. In addition, the solutions with small χ^2 values tend to be concentrated in the region of 2 [meV] $\leq m_{\beta\beta} \leq 4$ [meV].

Based on the definition in Eq. (2.18), the mass scale of left-handed neutrino determined by Eq. (2.2) does not guarantee that the experimental constraints on $\Sigma_i m_i$ and $m_{\beta\beta}$ are automatically satisfied. Therefore, the generated data is not necessarily within the experimental constraints due to selection bias. Given this fact, the tendencies described in the previous paragraph seem to emerge as the diffusion model learns some features to reproduce the




FIG. 4. The histograms show the distribution of the Majorana phase α_{31} and that absolute sine $|\sin \alpha_{31}|$, respectively. It turns out that few solutions are found near $\alpha_{31} = 0$ [deg]. The median of $|\sin \alpha_{31}|$ is 0.749, with the first and third quartiles at 0.415 and 0.907, respectively. These values suggest that CP symmetry is significantly broken.




FIG. 5. The distribution of the effective Majorana neutrino mass $m_{\beta\beta}$ with respect to the sum of neutrino masses in the left panel and the mass of electron neutrino in the right panel. The color bars show χ^2 values with $P_\ell = \{\Delta m_{21}^2, \Delta m_{31}^2, s_{12}^2, s_{23}^2, s_{13}^2\}$ in Eq. (2.22). The white area is the 95% CL allowed regions for the normal ordering, and the gray area separated by the red boundary means outside of those regions based on NuFIT 6.0 [29]. All the solutions generated by the diffusion model satisfy the restrictions from experiments along the boundaries.

experimental results. Note that one can arrive at almost the same distributions, even when we change the neural network architecture. Since it is expected that extracting the features from the neural network will deepen our understanding of the lepton sectors, one of our future tasks is to analyze the data with enhanced interpretability referring to explainable AI.

IV. CONCLUSIONS

From various perspectives, numerous efforts have been made to comprehend the flavor structure of the quarks and the leptons in a unified framework. In particular, the nature of neutrinos remains enigmatic in many respects. We have considered a simple model with the type I seesaw mechanism and attempted to generate mass matrices that

reproduce the experimental results through the diffusion model, known as a type of generative AI. The diffusion model, which uses real experimental values as conditional labels, provides a novel methodological approach for exploring the distribution and correlations of physical quantities within the space defined by the adopted labels. The results should therefore be interpreted in an exploratory sense, rather than as definitive physical predictions, and may serve as guidance for future experimental studies.

In this paper, we have constructed the diffusion model with CFG to explore the flavor structure of leptons. We assumed the type I seesaw mechanism and prepared various mass matrices for the lepton sector based on uniformly distributed random numbers, as detailed in Sec. II A. During the diffusion process described in Sec. II B 1, the neural network was trained to predict noise, which was

added to the initial data, using actual noisy data and conditional labels. In Sec. II B 2, we then generated new mass matrices of neutrinos utilizing real observables as conditional labels, and in Sec. II B 3 we applied the transfer learning to improve the consistency of the generated samples with experimental data. The tuned-network generated the 104 solutions that satisfy the 3σ constraints of neutrino masses and mixing angles. In Sec. III, we analyzed the characteristics of the generated results and found that the CP phases and the sums of neutrino masses, which were not included in the conditional labels, exhibited nontrivial tendencies under the requirement of reproducing the experimental observables. We also showed that the effective Majorana neutrino masses calculated from the generated data are concentrated near the boundaries of the currently allowed confidence intervals, providing qualitative insight into the structure of the obtained solutions.

Before closing our paper, we will mention possible future works:

- (i) In this study, a straightforward model with only the type I seesaw mechanism was employed to study the CP violation, the effective Majorana neutrino mass, etc. For future analyses, we can consider lepton flavor violation in theories beyond the Standard Model, allowing the exploration of decay branching ratios in light of existing experimental results. Since the diffusion-model-based analysis is independent of the specific details of the underlying model, it provides a flexible framework for extracting qualitative insights from existing observables.
- (ii) Within the scope of this study, we have not examined whether correlations exist between the values of the Yukawa matrices generated by the diffusion model. In particular, if there is any symmetry in the Yukawa matrix, it is derived as a favorable structure to reproduce existing experimental results. Traditionally, the analysis of flavor models often begins with the assumption of certain symmetries. Conversely, an inverse approach utilizing the diffusion model or other generative AI techniques is expected to facilitate the verification of flavor models from a different perspective.
- (iii) When we generate data using DDPM with a neural network trained on $T = 1000$, it is necessary to perform 1,000 denoising steps for each generated data G . For large-scale generation, reducing the number of steps is desirable to improve calculational speed. While DDPM is based on stochastic denoising, Ref. [36] proposes denoising diffusion implicit models (DDIMs), which employ a deterministic generation process. It has been reported that DDIM can achieve accurate image generation while decreasing the number of inference steps. In the previous work, both DDPM and DDIM have been implemented in a unified framework with a single parameter. Additionally, Ref. [37] describes an

approach to improve the calculational speed by integrating mechanisms from generative adversarial networks (GANs), another type of generative AI, into the diffusion models. Furthermore, Ref. [38] proposes a method called model-guidance, which has demonstrated accuracy surpassing that of the CFG. Although the diffusion models have achieved remarkable results in image generation, it remains uncertain whether enhanced methods like DDIM will be effective when the neural networks learn eigenvalue decomposition to obtain conditional labels with the hierarchical structure. Comparative studies are needed to confirm whether the latest developments in diffusion models are also valid for the analysis of the flavor models.

ACKNOWLEDGMENTS

This work was supported in part by Kyushu University's Innovator Fellowship Program (S. N.), JSPS KAKENHI Grants No. JP25KJ1927 (S. N.), No. JP23H04512 (H. O.), No. JP25H01539 (H. O.).

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX A: FORMULATION OF DIFFUSION MODELS

In this section, we introduce the denoising diffusion probabilistic model (DDPM), which is one of the formulations of the diffusion model by Ref. [21]. The DDPM consists of a diffusion process that gradually adds noise to input data and a reverse process that conversely removes the noise. This framework provides an intuitive definition of the diffusion models. When a variable is explicitly specified in our notation, $\mathcal{N}(x; \mu, \sigma)$ is a normal distribution with a random variable x , an average μ , and a variance σ . Furthermore, a conditional probability $P(A|B)$ denotes the probability of A given the condition B .

1. Diffusion process and reverse process

Consider the following Markov process that adds noise to an original input data x_0 to obtain a series of new data $x_1, x_2, \dots, x_T$

$$q(x_{1:T}|x_0) = \prod_{t=1}^T q(x_t|x_{t-1}), \quad (\text{A1})$$

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{\alpha_t}x_{t-1}, \beta_t), \quad (\text{A2})$$

with $x_{i:j} = x_i, x_{i+1}, \dots, x_j$. Here, the parameters $0 < \beta_1 < \beta_2 < \dots < \beta_T < 1$ decide the variance of $\mathcal{N}$, and α_t is defined as $\alpha_t = 1 - \beta_t$. Thus, $\{\alpha_1, \dots, \alpha_T, \beta_1, \dots, \beta_T\}$ are called noise schedules. By the definition of α and the magnitude of β , the amount of original input data becomes smaller and smaller, and noise grow to be dominant. In other words, $q(x_T|x_0) \simeq \mathcal{N}(x_T; 0, 1)$ holds for any x_0 , so the final data x_T is a complete noise. By using a normal distribution for the adding noise, the conditional probability at any time t under a condition x_0 can be written in the following analytical expression:

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, \bar{\beta}_t), \quad (\text{A3})$$

with

$$\bar{\alpha}_t = \prod_{s=1}^t \alpha_s, \quad \bar{\beta}_t = 1 - \bar{\alpha}_t. \quad (\text{A4})$$

The reproductive property of the normal distribution leads to this result.

Now, consider the complete noise $\mathcal{N}(x_T; 0, 1)$ as an initial value and follow the diffusion process in the reverse direction. This is called a reverse process. The conditional probability $p_\theta(x_{t-1}|x_t)$ at each step is estimated by a mathematical model characterized by parameters θ . In this study, a neural network was used as the mathematical model. The whole Markov process is described as follows:

$$p_\theta(x_{0:T}) = p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t), \quad (\text{A5})$$

$$p(x_T) = \mathcal{N}(x_T; 0, 1). \quad (\text{A6})$$

The probability distribution $p_\theta(x_{0:T})$ that generates a set of data $x_{0:T} = x_0, x_1, \dots, x_T$ is called a likelihood. If T is large and the increments of β are sufficiently small, the diffusion

process and the reverse process have the same functional form [39]. Therefore, $p_\theta(x_{t-1}|x_t)$ for each step follows the following equation:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \sigma_\theta(x_t, t)). \quad (\text{A7})$$

According to Ref. [40], a fixed variance $\sigma_\theta(x_t, t) = \sigma_t^2$, which is independent of the parameter θ , ensures stability and accuracy of learning:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \sigma_t^2). \quad (\text{A8})$$

It is known that the same results can be obtained in the reverse process whether $\sigma_t^2 = \beta_t$ or $\sigma_t^2 = \bar{\beta}_t$ is chosen.

2. Learning in diffusion process

When optimizing the parameters θ of the mathematical model with realistic computational costs, we often perform maximum likelihood estimation, which is formulated by maximizing the evidence lower bound of the log-likelihood. This is detailed below. The original input data x_0 is an observable variable, while the remaining $x_{1:T}$ is latent. Thus, in the reverse process, the likelihood $p_\theta(x_0)$ of the observed variable x_0 is obtained by integrating these latent variables as follows:

$$p_\theta(x_0) = \int dx_{1:T} p_\theta(x_{0:T}). \quad (\text{A9})$$

Since it is practically impossible to evaluate this integral, maximizing the likelihood itself is difficult. Therefore, we consider the following inequality that holds for the log-likelihood:

$$-\log p_\theta(x_0) \leq \mathbb{E}_{q(x_{1:T}|x_0)} \left[-\log \frac{p_\theta(x_{0:T})}{q(x_{1:T}|x_0)} \right] \quad (\text{A10})$$

$$= \mathbb{E}_{q(x_{1:T}|x_0)} \left[-\log \frac{p_\theta(x_0|x_1)p_\theta(x_1|x_2) \cdots p_\theta(x_{T-1}|x_T)p_\theta(x_T)}{q(x_T|x_{T-1})q(x_{T-1}|x_{T-2}) \cdots q(x_1|x_0)} \right] \quad (\text{A11})$$

$$= \mathbb{E}_{q(x_{1:T}|x_0)} \left[-\log p_\theta(x_T) - \sum_{t=1}^T \log \frac{p_\theta(x_{t-1}|x_t)}{q(x_t|x_{t-1})} \right] \equiv L(\theta). \quad (\text{A12})$$

In the transformation from Eq. (A10) to Eq. (A11), p_θ and q are decomposed into products based on the Markov property of diffusion and reverse processes. In the end, maximum likelihood estimation for parameter updating is achieved by minimizing $L(\theta)$.

By rewriting $L(\theta)$ using Kullback-Leibler divergence⁵

$$D_{\text{KL}}(P||Q) = \int dx P(x) \log \frac{P(x)}{Q(x)}, \quad (\text{A13})$$

we obtain the following representation [40]:

⁵ $D_{\text{KL}}(P||Q)$ is a quantified evaluation of the difference between the two probability distributions P, Q .

$$L(\theta) = \mathbb{E}_q \left[\underbrace{D_{\text{KL}}(q(x_T|x_0)||p(x_T))}_{L_T} + \sum_{t>1} \underbrace{D_{\text{KL}}(q(x_{t-1}|x_t, x_0)||p_\theta(x_{t-1}|x_t))}_{L_{t-1}} - \underbrace{\log p_\theta(x_0|x_1)}_{L_0} \right]. \quad (\text{A14})$$

L_T is not related to the minimization of $L(\theta)$ because it does not include the parameter θ . Regarding L_0 , the last reverse process is considered a normal distribution that is independent for each dimension of the input data. Then, L_0 is evaluated by discretizing the interval $[-1, 1]$ into k pieces as follows:

$$p_\theta(x_0|x_1) = \prod_{i=1}^d \int_{\sigma_-(x_0^i)}^{\sigma_+(x_0^i)} dx \mathcal{N}(x; \mu_\theta^i(x_1, 1), \sigma_1^2) \quad (\text{A15})$$

$$\sigma_+(x) = x + \frac{1}{k}, \quad \sigma_-(x) = x - \frac{1}{k}. \quad (\text{A16})$$

L_{t-1} is evaluated as follows using a constant C that is independent of θ and a noise ϵ following a standard normal distribution $\mathcal{N}(0, 1)$:

$$L_{t-1} - C = \mathbb{E}_{x_0, \epsilon} \left[\frac{1}{2\sigma_t^2} \left\| \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{\beta_t}} \epsilon \right) - \mu_\theta(x_t, t) \right\|^2 \right], \quad (\text{A17})$$

$$x_t = \sqrt{\alpha_t} x_0 + \sqrt{\beta_t} \epsilon. \quad (\text{A18})$$

To simplify the representation, a variable transformation is applied as

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{\beta_t}} \epsilon_\theta \right). \quad (\text{A19})$$

Here, ϵ_θ is predicted noise from the mathematical model. With this transformation, the right-hand side of Eq. (A17) is summarized in the following form:

$$L_{t-1} - C = \mathbb{E}_{x_0, \epsilon} \left[\frac{\beta_t^2}{2\sigma_t^2 \alpha_t \beta_t} \|\epsilon - \epsilon_\theta(x_t, t)\|^2 \right], \quad (\text{A20})$$

A rigorous evaluation of $L(\theta)$ is calculated by the sum of L_T , L_{t-1} , and L_0 . In practical optimization, it is known that the following simplified function is sufficient [40]:

$$L_{\text{simple}}(\theta) = \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2], \quad (\text{A21})$$

$$x_t = \sqrt{\alpha_t} x_0 + \sqrt{\beta_t} \epsilon. \quad (\text{A22})$$

In summary, the learning of DDPM proceeds in the following three steps. First, x_t is constructed by adding noise to the initial data x_0 based on the noise schedules. Second, the noise ϵ added to the noisy data x_t is estimated with the mathematical model. Finally, the parameters θ of the mathematical model are updated so that the difference between the actual noise ϵ and the predicted noise ϵ_θ is reduced. These procedures are repeated until L_{simple} converges.

3. Data generation in reverse process

During generating data through the reverse process, the perfect noise x_T is used as initial data, and x_{t-1} based on x_t is sampled in sequence according to Eq. (A8). This sampling method is called ancestral sampling. The average of the normal distribution is $\mu_\theta(x_t, t)$, which was estimated in Eq. (A19). Thus, sampling from $p_\theta(x_{t-1}|x_t)$ corresponds to determining x_{t-1} by the following equation:

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{\beta_t}} \epsilon_\theta \right) + \sigma_t u_t, \quad (\text{A23})$$

where u_t is a disturbance following a standard normal distribution.

In summary, data generation by the reverse process of DDPM proceeds in the following two steps. First, the complete noise x_T is used as initial data, and noise ϵ_θ is predicted for each time using the mathematical model already trained in the diffusion process. Second, ϵ_θ is removed from the data x_t and some disturbance u_t is added. The final result x_0 is obtained as new data by repeating these procedures.

APPENDIX B: CONDITIONAL DIFFUSION MODELS

The success of the diffusion models is attributed not only to its ability to generate diverse data from complete noise but also to its capacity to produce preferred outputs based on arbitrary conditional labels. For instance, the diffusion models can generate images that correspond to any given text. Additionally, it can create high-resolution images from low-resolution inputs, and this process is known as super-resolution. To ensure smooth generation under various conditions, it is preferable to utilize a single trained model that can manage multiple labels, rather than retraining the model for each individual label.

In this section, we provide a brief introduction to classifier guidance (CG) [41], which enables conditional generation, and classifier-free guidance (CFG) [42], which further enhances the learning process. The following explanation is based on a formulation known as the score-based models (SBMs) [43,44]. Although the derivations of SBM and DDPM are different, the final evaluation functions of both models are similar [Eq. (A21)]. It is established that optimal solutions for minimizing the evaluation functions are also consistent, and in most cases, both definitions can be applied

as diffusion models [45]. Therefore, the derivation of SBM is not detailed in this paper.

1. Classifier guidance

In a well-trained diffusion model, the gradient of the log-likelihood in the diffusion process is calculated as follows:

$$\begin{aligned} s(x_t, t) &= \nabla_{x_t} \log q(x_t) \\ &= -\frac{\epsilon_\theta(x_t, t)}{\sigma_t}. \end{aligned} \quad (\text{B1})$$

This gradient s is called a score. On the other hand, in general, conditional probability can be transformed based on Bayes' theorem

$$q(x_t|c) = \frac{q(c|x_t)q(x_t)}{q(c)}. \quad (\text{B2})$$

Thus, given a label c for a diffusion model, its conditional score is calculated as follows:

$$\begin{aligned} s(x_t, t, c) &= \nabla_{x_t} \log q(x_t|c) \\ &= \nabla_{x_t} \log q(c|x_t) + \nabla_{x_t} \log q(x_t) \\ &= -\frac{\epsilon_\theta(x_t, t)}{\sigma_t} + \nabla_{x_t} \log q(c|x_t) \\ &= -\frac{\epsilon_\theta(x_t, t) - \sigma_t \nabla_{x_t} \log q(c|x_t)}{\sigma_t}. \end{aligned} \quad (\text{B3})$$

In CG, $\nabla_{x_t} \log q(c|x_t)$ is estimated using a learned classification model $p_\theta(c|x_t)$. In other words, the new classification model (not the diffusion model) induces the generation of data that fits designated labels. Specifically, ϵ_θ is corrected as follows:

$$\begin{aligned} \hat{\epsilon}_\theta(x_t, t, c) &= \epsilon_\theta(x_t, t) - \gamma \sigma_t \nabla_{x_t} \log q(c|x_t) \\ &= \epsilon_\theta(x_t, t) - \gamma \sigma_t \nabla_{x_t} \log p_\theta(c|x_t), \end{aligned} \quad (\text{B4})$$

and the reverse process depends on the corrected noise $\hat{\epsilon}_\theta$ instead of the simple noise ϵ_θ to generate data that matches the label c . $\gamma > 0$ is called a guidance scale. $\gamma = 1$ matches the original conditional probability, and $\gamma > 1$ emphasizes the conditional label.

There are two practical issues with CG. The first is that it requires the score $\nabla_{x_t} \log p_\theta(c|x_t)$, which depends on time t . To estimate this score, a classification model $p_\theta(c|x_t)$ must be trained for each various noise level, so a computational cost is high. The second is that poor quality of conditional generation is realized when data x and labels c have little or no relation to each other.

2. Classifier-free guidance

CFG does not use the classification model $p_\theta(c|x_t)$, but only the diffusion model to learn conditional scores directly. Specifically, a diffusion model that can receive the label c is

constructed, and the corrected noise $\hat{\epsilon}_\theta$ is transformed as follows:

$$\begin{aligned} \hat{\epsilon}_\theta(x_t, t, c) &= \epsilon_\theta(x_t, t) - \gamma \sigma_t \nabla_{x_t} \log q(c|x_t) \\ &= \epsilon_\theta(x_t, t) - \gamma \sigma_t \nabla_{x_t} \log \frac{q(x_t|c)q(c)}{q(x_t)} \\ &= \epsilon_\theta(x_t, t) \\ &\quad - \gamma [\sigma_t \nabla_{x_t} \log q(x_t|c) - \sigma_t \nabla_{x_t} \log q(x_t)] \\ &= (1 - \gamma) \epsilon_\theta(x_t, t) + \gamma \epsilon_\theta(x_t, t, c). \end{aligned} \quad (\text{B5})$$

In other words, the corrected noise $\hat{\epsilon}_\theta$ is determined by the linear combination of the unlabeled noise and the labeled noise. The coefficient γ is called the CFG scale, which was adopted as 8.0 in our analysis. The larger the CFG scale, the more faithful the data are to the label c , but the diversity of generated results tends to be lost. On the other hand, when $\gamma < 1$, it emphasizes the diversity.

At first glance, the evaluation of Eq. (B5) require two mathematical models, $\epsilon_\theta(x_t, t, c)$ and $\epsilon_\theta(x_t, t)$. However, a common mathematical model can be used by using $\emptyset$ to denote the absence of conditional labels as $\epsilon_\theta(x_t, t) = \epsilon_\theta(x_t, t, c = \emptyset)$. For $\emptyset$, a learned embedding vector or a zero vector $\emptyset = 0$ is often used. In updating the parameter θ , $c = \emptyset$ is adopted with low probability (10%–20%), and the learning of CFG proceeds with mixing the cases with and without labels.

CFG does not require an independent classification model. It needs just dropping out the conditional part with a certain probability in usual learning. Therefore, the learning process with CFG can be significantly simplified in comparison to CG. Furthermore, the correspondence between data x and labels c can be predicted with good accuracy, so high-quality conditional generation can be realized.

APPENDIX C: TRANSFER LEARNING

In this section, we formulate the transfer learning. For reference, Ref. [46] provides an overview of transfer learning. The domain for preparing a learned neural network is referred to as the source domain, while the domain in which the same network is applied is called the target domain. When constructing a neural network Y_s using input data X_s from the source domain, a multilayer structure of the network is expressed as follows based on Eq. (2.11).

$$Y_s = X_P \circ X_{P-1} \circ \cdots \circ X_1(X_s), \quad (\text{C1})$$

with a whole number of the layers P . Then, we focus on a submodel $\Phi = X_Q \circ X_{Q-1} \circ \cdots \circ X_1$ up to the Q layers of the trained model Y_s , where $Q < P$. In a well-trained neural network in the source domain, parameters of Φ include features that are useful for predicting Y_s .

If there are common features between the source and target domains, Φ is expected to be useful for prediction in the target domain also. Thus, under input data X_t from the

target domain, training a model of the form $Y_t = f_t \circ \Phi(X_t)$ is called transfer learning through feature extraction. Here, f_t represents an arbitrary neural network. Since Φ has already been trained on the source domain, its parameters remain fixed. In other words, since learning in the target domain is limited to f_t , there is no need to update all of the parameters of Y_t . Furthermore, transfer learning enables the construction of a neural network with high accuracy even with a limited amount of data X_t .

The fine-tuning used in this study is a type of transfer learning. In this approach, $Q = 0$ is adopted in the aforementioned formulation, and the parameters of the network Y_s trained in the source domain are set as the initial values of Y_t . Then, Y_t is retrained using the dataset X_t from the target domain. While allowing all parameters to be updated ensures flexibility of learning in the target domain, a large amount of data X_t should be prepared compared to the transfer learning with $Q > 0$.

-
- [1] C.D. Froggatt and H.B. Nielsen, Hierarchy of quark masses, Cabibbo angles and CP violation, *Nucl. Phys.* **B147**, 277 (1979).
 - [2] G. Altarelli and F. Feruglio, Discrete flavor symmetries and models of neutrino mixing, *Rev. Mod. Phys.* **82**, 2701 (2010).
 - [3] H. Ishimori, T. Kobayashi, H. Ohki, Y. Shimizu, H. Okada, and M. Tanimoto, Non-Abelian discrete symmetries in particle physics, *Prog. Theor. Phys. Suppl.* **183**, 1 (2010).
 - [4] D. Hernandez and A.Y. Smirnov, Lepton mixing and discrete symmetries, *Phys. Rev. D* **86**, 053014 (2012).
 - [5] S.F. King and C. Luhn, Neutrino mass and mixing with discrete symmetry, *Rep. Prog. Phys.* **76**, 056201 (2013).
 - [6] S.F. King, A. Merle, S. Morisi, Y. Shimizu, and M. Tanimoto, Neutrino mass and mixing: From theory to experiment, *New J. Phys.* **16**, 045018 (2014).
 - [7] S. T. Petcov, Discrete flavour symmetries, neutrino mixing and leptonic CP violation, *Eur. Phys. J. C* **78**, 709 (2018).
 - [8] T. Kobayashi, H. Ohki, H. Okada, Y. Shimizu, and M. Tanimoto, *An Introduction to Non-Abelian Discrete Symmetries for Particle Physicists* (Springer, Berlin, Heidelberg, 2022), 10.1007/978-3-662-64679-3.
 - [9] F. Feruglio, Are neutrino masses modular forms?, in *From My Vast Repertoire: Guido Altarelli's Legacy*, edited by A. Levy, S. Forte, and G. Ridolfi (2019), pp. 227–266, 10.1142/9789813238053_0012.
 - [10] T. Kobayashi and M. Tanimoto, Modular flavor symmetric models, *Int. J. Mod. Phys. A* **39**, 2441012 (2024).
 - [11] G.-J. Ding and S. F. King, Neutrino mass and mixing with modular symmetry, *Rep. Prog. Phys.* **87**, 084201 (2024).
 - [12] T. Kobayashi and H. Otsuka, Non-invertible flavor symmetries in magnetized extra dimensions, *J. High Energy Phys.* **11** (2024) 120.
 - [13] T. Kobayashi, H. Otsuka, and M. Tanimoto, Yukawa textures from non-invertible symmetries, *J. High Energy Phys.* **12** (2024) 117.
 - [14] J. Hisano, T. Moroi, K. Tobe, and M. Yamaguchi, Lepton flavor violation via right-handed neutrino Yukawa couplings in supersymmetric standard model, *Phys. Rev. D* **53**, 2442 (1996).
 - [15] J. Hisano and D. Nomura, Solar and atmospheric neutrino oscillations and lepton flavor violation in supersymmetric models with the right-handed neutrinos, *Phys. Rev. D* **59**, 116005 (1999).
 - [16] J.A. Casas and A. Ibarra, Oscillating neutrinos and $\mu \rightarrow e, \gamma$, *Nucl. Phys.* **B618**, 171 (2001).
 - [17] D. Falcone, Inverting the seesaw formula, *Phys. Rev. D* **68**, 033002 (2003).
 - [18] S. Nishimura, C. Miyao, and H. Otsuka, Exploring the flavor structure of quarks and leptons with reinforcement learning, *J. High Energy Phys.* **12** (2023) 021.
 - [19] S. Nishimura, C. Miyao, and H. Otsuka, Reinforcement learning-based statistical search strategy for an axion model from flavor, *J. High Energy Phys.* **10** (2025) 043.
 - [20] P. Devlin, J.-W. Qiu, F. Ringer, and N. Sato, Diffusion model approach to simulating electron-proton scattering events, *Phys. Rev. D* **110**, 016030 (2024).
 - [21] J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, and S. Ganguli, Deep unsupervised learning using nonequilibrium thermodynamics, *arXiv:1503.03585*.
 - [22] P. Minkowski, $\mu \rightarrow e\gamma$ at a rate of one out of 10^9 muon decays?, *Phys. Lett.* **67B**, 421 (1977).
 - [23] T. Yanagida, Horizontal gauge symmetry and masses of neutrinos, *Conf. Proc. C* **7902131**, 95 (1979), <https://inspirehep.net/files/19a4de25a11cc3ea4cc4e0bd0f45bdbc>.
 - [24] M. Gell-Mann, P. Ramond, and R. Slansky, Complex spinors and unified theories, *Conf. Proc. C* **790927**, 315 (1979), http://tuvalu.santafe.edu/~mgm/Site/Publications_files/MGM%2086.pdf.
 - [25] R. N. Mohapatra and G. Senjanovic, Neutrino mass and spontaneous parity nonconservation, *Phys. Rev. Lett.* **44**, 912 (1980).
 - [26] L. J. Hall, H. Murayama, and N. Weiner, Neutrino mass anarchy, *Phys. Rev. Lett.* **84**, 2572 (2000).
 - [27] Planck Collaboration, Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
 - [28] DESI Collaboration, DESI 2024 VI: Cosmological constraints from the measurements of baryon acoustic oscillations, *J. Cosmol. Astropart. Phys.* **02** (2025) 021.
 - [29] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, J. a. P. Pinheiro, and T. Schwetz, NuFit-6.0: Updated global analysis of three-flavor neutrino oscillations, *J. High Energy Phys.* **12** (2024) 216.
 - [30] KamLAND-Zen Collaboration, Search for the Majorana nature of neutrinos in the inverted mass ordering region with KamLAND-Zen, *Phys. Rev. Lett.* **130**, 051801 (2023).

- [31] Katrin Collaboration, Direct neutrino-mass measurement based on 259 days of KATRIN data, *Science* **388**, 180 (2025).
- [32] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, Flow matching for generative modeling, [arXiv:2210.02747](#).
- [33] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps, [arXiv:2206.00927](#).
- [34] H. Shen, J. Zhang, B. Xiong, R. Hu, S. Chen, Z. Wan *et al.*, Efficient diffusion models: A survey, [arXiv:2502.06805](#).
- [35] A. Nichol and P. Dhariwal, Improved denoising diffusion probabilistic models, [arXiv:2102.09672](#).
- [36] J. Song, C. Meng, and S. Ermon, Denoising diffusion implicit models, [arXiv:2010.02502](#).
- [37] L. T. Trinh and T. Hamagami, Latent denoising diffusion gan: Faster sampling, higher image quality, *IEEE Access* **12**, 78161 (2024).
- [38] Z. Tang, J. Bao, D. Chen, and B. Guo, Diffusion models without classifier-free guidance, [arXiv:2502.12154](#).
- [39] W. Feller, On the theory of stochastic processes, with particular reference to applications, in *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability* (1949), pp. 403–432.
- [40] J. Ho, A. Jain, and P. Abbeel, Denoising diffusion probabilistic models, [arXiv:2006.11239](#).
- [41] P. Dhariwal and A. Nichol, Diffusion models beat gans on image synthesis, [arXiv:2105.05233](#).
- [42] J. Ho and T. Salimans, Classifier-free diffusion guidance, [arXiv:2207.12598](#).
- [43] Y. Song and S. Ermon, Generative modeling by estimating gradients of the data distribution, [arXiv:1907.05600](#).
- [44] Y. Song and S. Ermon, Improved techniques for training score-based generative models, [arXiv:2006.09011](#).
- [45] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, Score-based generative modeling through stochastic differential equations, [arXiv:2011.13456](#).
- [46] H. Asmaul, M. Ethel, G. Jigme, A. Zulfikar, A. Zeyar, and A. A. Mohammad, Transfer learning: A friendly introduction, *J. Big Data* **9**, 102 (2022).