Vector Kuramoto Diffusion on Hyperspheres for Phase-Based EEG Synchronization Modeling

Vivaan Daxini
Archbishop Mitty High School
San Jose, CA, USA
vivdax54@gmail.com

***Abstract*—We present a Kuramoto-based diffusion framework for modeling oscillator-valued data on non-Euclidean state spaces. Standard diffusion models are usually formulated in Euclidean space and do not naturally preserve the geometry of phase or directional variables. To address this limitation, we formulate a vector Kuramoto diffusion process on the hypersphere, where oscillator states are constrained to remain on S^{d-1} and synchronization corresponds to tangent-space alignment. The classical circular phase model is recovered as the S^1 special case.**

The empirical study focuses on this circular specialization, which is appropriate for phase-valued data. First, image data from the SVHN dataset are converted into orientation fields, where each pixel is represented as a phase. Second, neural phase data are extracted from EEG recordings of healthy controls and patients with Alzheimer's disease and frontotemporal dementia. The model is trained for 1500 epochs and achieves a reconstruction Fréchet Inception Distance (FID) score of 54.2 on SVHN. Generated samples preserve recognizable visual structure and diversity, while phase distributions and Kuramoto synchronization statistics follow those of real data. In EEG analysis, the model reconstructs phase structure across clinical groups and provides a generative description of neural synchronization patterns. These results demonstrate that Kuramoto-based diffusion can model both local phase statistics and global synchronization structure in the circular setting, while the hyperspherical formulation provides a natural extension toward higher-dimensional directional data.

***Index Terms*—diffusion models, Kuramoto model, phase synchronization, score-based generative models, EEG, Alzheimer's disease, frontotemporal dementia, circular manifolds**

I. Introduction

A natural theoretical framework for studying synchronization phenomena is provided by the Kuramoto model, a canonical dynamical system describing the collective behavior of coupled oscillators [1]. In its classical form, each oscillator is characterized by a phase variable on the circle, and pairwise sinusoidal interactions promote phase alignment between oscillators. Despite its simplicity, the Kuramoto model captures a rich range of emergent phenomena including synchronization phase transitions, chimera states, and metastable dynamics [2]. Because neural oscillations are often analyzed in terms of phase relationships between signals, the Kuramoto framework has become an important tool for modeling large-scale brain synchronization [3].

Recent advances in generative modeling have led to the development of diffusion models, which have achieved state-of-the-art performance in tasks such as image synthesis, audio generation, and molecular design [4]. Diffusion models learn complex data distributions by defining a stochastic process that gradually corrupts data with noise together with a learned reverse-time process that reconstructs samples from noise. These models are typically formulated in Euclidean spaces and rely on Gaussian diffusion processes, which are well suited for vector-valued data such as images.

However, many real-world datasets possess intrinsic geometric structure that is not naturally represented in Euclidean space. In particular, phase variables arising from oscillatory systems lie on the circle, forming a periodic manifold. More generally, oscillator states may be represented as unit vectors constrained to hyperspheres, where synchronization corresponds to alignment on a curved state space rather than translation in a Euclidean vector space. Standard Gaussian diffusion processes do not respect this geometry and may therefore fail to capture the underlying structure of phase-based or directional data. This limitation is particularly relevant for neural recordings, where the dynamics of interest are often expressed in terms of phase synchronization across oscillatory channels.

In this work we introduce a Kuramoto-based diffusion framework for structured oscillator-valued data. We first formulate the model in a vector form on the hypersphere, where each oscillator state is a unit vector and the interaction dynamics are constrained to the tangent space of the sphere. The classical phase Kuramoto model on the circle is then recovered as the two-dimensional special case. This hyperspherical formulation provides a natural way to extend score-based generative modeling beyond Euclidean spaces while preserving the geometry of directional and phase-valued variables.

For the empirical part of this study, we focus on the circular special case motivated by phase-valued data. Our approach integrates stochastic Kuramoto dynamics into the forward diffusion process, replacing the linear drift typically used in conventional diffusion models with interactions that promote phase alignment between oscillators. By combining synchronization dynamics with stochastic diffusion, the resulting model naturally operates on circular phase manifolds while preserving the essential structure of oscillatory systems.

We apply this framework to two complementary settings. First, we demonstrate the approach on image data by rep-

The author is with Archbishop Mitty High School, San Jose, CA, USA.

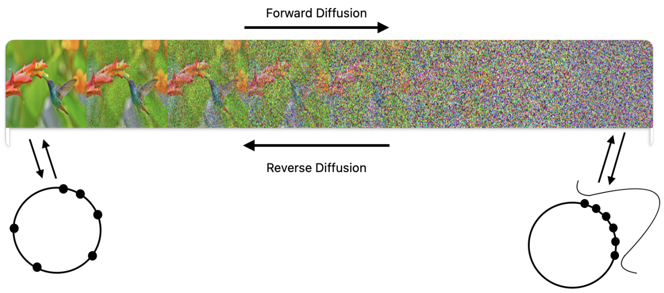

Fig. 1. The forward diffusion process (top, left to right) progressively corrupts a structured image into noise. The reverse diffusion process (top, right to left) reconstructs structured samples from noise using a learned score function. Unlike standard diffusion models operating in Euclidean space, the latent representation is defined on a circular phase (bottom), where each pixel corresponds to an oscillator with a phase. The forward process corresponds to stochastic Kuramoto dynamics that randomize phases, while the reverse process drives the system toward synchronized phase configurations that reconstruct coherent images.

resenting local image orientations as phase variables and learning generative models of orientation fields using the SVHN dataset [5]. Second, we apply the model to neural phase data extracted from EEG recordings in an OpenNeuro dataset containing subjects with Alzheimer's disease, frontotemporal dementia, and healthy controls [7]. In this setting, the Kuramoto diffusion model provides a generative description of neural synchronization patterns and allows us to investigate how disease alters the distribution of collective phase dynamics.

Our main contribution is the formulation of a diffusion model whose dynamics are governed by a stochastic Kuramoto system on non-Euclidean state spaces. The proposed hyperspherical formulation generalizes the standard circular phase model to vector-valued oscillator states, while the experiments in this paper validate the circular specialization on image orientation fields and EEG phase dynamics. Beyond its applications to neuroscience, the framework offers a general approach for modeling structured data defined on circular manifolds, hyperspheres, and related non-Euclidean domains.

II. METHODS

We integrate stochastic Kuramoto dynamics into a diffusion modeling pipeline. We first summarize score-based generative modeling, then formulate a vector Kuramoto diffusion process on the hypersphere. The standard phase model on the circle is recovered as the two-dimensional special case used in our image and EEG experiments.

A. Background: Score-Based Generative Models

Score-based generative models learn a data distribution $p(\mathbf{x})$ by defining a forward stochastic process that corrupts data with noise and a learned reverse process that maps noise back to structured samples. The central quantity is the score function $\nabla_{\mathbf{x}} \log p(\mathbf{x})$, which is approximated by a neural network $s_\theta(\mathbf{x}_t, t)$.

The forward diffusion process is commonly written as

$$d\mathbf{x}_t = f(\mathbf{x}_t, t)\, dt + g(t)\, d\mathbf{w}_t, \tag{1}$$

where f is the drift term, $g(t)$ is the diffusion coefficient, and $\mathbf{w}_t$ is a Wiener process. The corresponding reverse-time SDE is

$$d\mathbf{x}_t = \Big[f(\mathbf{x}_t, t) - g(t)^2 \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)\Big] dt + g(t)\, d\bar{\mathbf{w}}_t. \tag{2}$$

Since the true score is unknown, it is replaced by $s_\theta(\mathbf{x}_t, t)$ and learned using denoising score matching,

$$\mathcal{L}(\theta) = \mathbb{E}\left[\|s_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t \mid \mathbf{x}_0)\|^2\right]. \tag{3}$$

In standard Euclidean diffusion models, the forward process is usually Gaussian. In this work, we replace the Euclidean drift with Kuramoto-type synchronization dynamics defined on non-Euclidean manifolds.

B. Vector Kuramoto Diffusion on the Hypersphere

The classical Kuramoto model represents each oscillator by a scalar phase on the circle. To generalize this idea, we allow each oscillator to be represented by a unit vector $\mathbf{x}_i(t) \in S^{d-1} \subset \mathbb{R}^d$, with $\|\mathbf{x}_i(t)\| = 1$. In this setting, synchronization corresponds to alignment of vectors on the hypersphere rather than phase alignment on the circle. Such higher-dimensional Kuramoto-type models have been studied as natural extensions of phase synchronization to orientation and vector-valued states [11], [12].

Since the state must remain on the hypersphere, the velocity of each oscillator must lie in the tangent space of S^{d-1} at $\mathbf{x}_i$. This is enforced by the orthogonal projection

$$P^{\perp}_{\mathbf{x}_i} = \mathbf{I} - \mathbf{x}_i \mathbf{x}_i^{\top}. \tag{4}$$

For any vector $\mathbf{v} \in \mathbb{R}^d$, $P^{\perp}_{\mathbf{x}_i}\mathbf{v}$ removes the radial component of $\mathbf{v}$ and keeps only the component tangent to the sphere. Thus, the resulting dynamics change the direction of $\mathbf{x}_i$ without changing its norm.

A deterministic vector Kuramoto interaction can then be written as

$$\frac{d\mathbf{x}_i}{dt} = P^{\perp}_{\mathbf{x}_i}\left[\frac{K}{N}\sum_{j=1}^{N} \mathbf{x}_j\right], \tag{5}$$

where $K > 0$ controls the strength of alignment. The mean vector $\frac{1}{N}\sum_j \mathbf{x}_j$ points toward the average direction of the oscillator population. Projecting this mean direction onto the tangent space causes each oscillator to rotate along the sphere toward the population mean.

To define a diffusion process on the hypersphere, we add stochastic forcing and a reference direction $\boldsymbol{\mu}_{\mathrm{ref}} \in S^{d-1}$. The reference direction plays the role of a terminal attractor, analogous to the reference phase used in circular Kuramoto diffusion. The forward process is

$$d\mathbf{x}_i = P^{\perp}_{\mathbf{x}_i}\left[\frac{K(t)}{N}\sum_{j=1}^{N} \mathbf{x}_j + K_{\mathrm{ref}}(t)\boldsymbol{\mu}_{\mathrm{ref}}\right] dt + \sqrt{2D(t)}\, P^{\perp}_{\mathbf{x}_i} d\mathbf{W}_i - D(t)(d-1)\mathbf{x}_i\, dt. \tag{6}$$

Here $D(t)$ controls the noise strength, $K(t)$ controls the coupling between oscillators, and $K_{\mathrm{ref}}(t)$ controls attraction toward the reference direction. The noise term is projected onto the tangent space so that random perturbations also remain on the sphere. The final drift term, $-D(t)(d-1)\mathbf{x}_i\,dt$, is the Itô correction for Brownian motion on S^{d-1} [14]. This correction prevents the stochastic process from drifting away from the unit sphere in Itô form.

With a sufficiently strong reference term, the terminal distribution is concentrated around $\boldsymbol{\mu}_{\mathrm{ref}}$ and approaches a von Mises–Fisher distribution,

$$p(\mathbf{x}) = C_d(\kappa)\exp\bigl(\kappa\,\boldsymbol{\mu}_{\mathrm{ref}}^{\top}\mathbf{x}\bigr)\,, \qquad \mathbf{x}\in S^{d-1}, \tag{7}$$

where $\kappa \geq 0$ is a concentration parameter and $C_d(\kappa)$ is the normalizing constant. The von Mises–Fisher distribution is the standard analogue of a Gaussian-like concentrated distribution on the hypersphere [13]. Large κ concentrates probability mass near $\boldsymbol{\mu}_{\mathrm{ref}}$, while $\kappa = 0$ gives the uniform distribution on the sphere.

A natural synchronization statistic in this setting is the vector order parameter

$$R(t) = \left\| \frac{1}{N}\sum_{j=1}^{N}\mathbf{x}_j(t) \right\|. \tag{8}$$

When the oscillator directions are widely dispersed on the sphere, their mean vector cancels and $R(t) \approx 0$. When the oscillators are strongly aligned, the mean vector has large magnitude and $R(t) \approx 1$. Thus, $R(t)$ generalizes the classical Kuramoto order parameter from circular phase synchronization to vector synchronization on S^{d-1}.

C. Circular Special Case

The classical phase Kuramoto model is recovered when $d = 2$ by writing each unit vector as

$$\mathbf{x}_i(t) = \begin{bmatrix}\cos\theta_i(t)\\ \sin\theta_i(t)\end{bmatrix} \in S^1. \tag{9}$$

In this case, the tangent-projected vector dynamics reduce to phase-based Kuramoto interactions. The stochastic circular diffusion used in our experiments is

$$\begin{aligned} d\theta_i(t) = {} & \frac{K(t)}{N}\sum_{j=1}^{N}\sin\bigl(\theta_j(t)-\theta_i(t)\bigr)\,dt \\ & + K_{\mathrm{ref}}(t)\sin\bigl(\psi_{\mathrm{ref}}-\theta_i(t)\bigr)\,dt + \sqrt{2D(t)}\,dW_i(t). \end{aligned} \tag{10}$$

The reference phase ψ_{ref} defines the circular terminal distribution,

$$p(\theta) \propto \exp\bigl(\kappa\cos(\theta-\psi_{\mathrm{ref}})\bigr), \tag{11}$$

which is the von Mises distribution. The usual Kuramoto order parameter is

$$r(t)e^{i\psi(t)} = \frac{1}{N}\sum_{j=1}^{N} e^{i\theta_j(t)}. \tag{12}$$

D. Reverse Process and Training

The reverse-time generative process uses a learned score function $s_\theta(\boldsymbol{\theta}_t, t) \approx \nabla_{\boldsymbol{\theta}_t}\log p(\boldsymbol{\theta}_t)$:

$$\begin{aligned} d\theta_i(t) = \Biggl[& \frac{K(t)}{N}\sum_{j}\sin\bigl(\theta_j-\theta_i\bigr) + K_{\mathrm{ref}}(t)\sin\bigl(\psi_{\mathrm{ref}}-\theta_i\bigr) \\ & - g(t)^2 s_\theta(\boldsymbol{\theta}_t, t)\Biggr] dt + g(t)\,d\bar{W}_i(t). \end{aligned} \tag{13}$$

Each phase is embedded using trigonometric features,

$$\mathrm{embed}(\theta) = [\sin\theta, \cos\theta], \tag{14}$$

so that the score network receives a representation consistent with circular geometry.

Training uses denoising score matching with a wrapped transition density. For a deterministic drift $f(\theta, t)$ defined by (10), the circular transition is approximated by

$$\begin{aligned} p(\theta_t \mid \theta_{t-1}) \approx & \sum_{k=-K}^{K}\frac{1}{\sqrt{4\pi D(t-1)}} \\ & \times \exp\left[-\frac{\bigl(\theta_t-\theta_{t-1}-f(\theta_{t-1}, t-1)+2\pi k\bigr)^2}{4D(t-1)}\right], \end{aligned} \tag{15}$$

and the loss is

$$\mathcal{L} = \mathbb{E}\left[\bigl\|s_\theta(\theta_t, t) - \nabla_{\theta_t}\log p(\theta_t \mid \theta_{t-1})\bigr\|^2\right]. \tag{16}$$

E. Adaptation to Image Data

We validate the circular specialization on the SVHN dataset [5], [6], which contains color images of street-view house numbers. Since the model operates on angular variables, each image is mapped to a local orientation field. For image intensity $I(x, y)$, we compute

$$G_x = \frac{\partial I}{\partial x}, \qquad G_y = \frac{\partial I}{\partial y}, \tag{17}$$

and define the phase at each pixel as

$$\theta(x, y) = \arctan\left(\frac{G_y}{G_x}\right) \in [-\pi, \pi]. \tag{18}$$

The resulting orientation field is flattened into an oscillator state vector $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$, with $N = 32\times 32$, and evolved using the circular Kuramoto diffusion process.

F. Adaptation to Neural Phase Data

We analyze resting-state EEG recordings from OpenNeuro dataset `ds004504` [7], containing healthy controls, Alzheimer's disease patients, and frontotemporal dementia patients. For each EEG channel i, the analytic signal is computed using the Hilbert transform,

$$z_i(t) = x_i(t) + i\,\mathcal{H}[x_i(t)], \tag{19}$$

and the instantaneous phase is

$$\theta_i(t) = \arg\bigl(z_i(t)\bigr). \tag{20}$$

Fig. 2. Comparison between real SVHN images (left) and samples generated by the Kuramoto diffusion model (right). The generated images capture the structure and variability of street number data, demonstrating the model's ability to learn digit-like patterns from natural scenes using a phase-based representation.

This produces multichannel phase vectors $\boldsymbol{\theta}(t) = (\theta_1(t), \ldots, \theta_N(t))$. We focus on the alpha band, 8–13 Hz, which is commonly used in studies of dementia-related oscillatory changes [8].

Each EEG phase vector is interpreted as the state of N coupled oscillators on S^1. Thus, the EEG experiments validate the circular special case of the proposed hyperspherical framework. After training, the reverse process generates synthetic phase configurations that follow the learned distribution of neural synchronization patterns.

III. Results

We evaluate the proposed framework in the circular setting, which is the S^1 specialization of the hyperspherical vector Kuramoto diffusion model. This setting is appropriate for both image orientation fields and EEG phase data, since both are naturally represented by angular variables. We first evaluate the model on the SVHN dataset by comparing generated samples, phase statistics, and synchronization properties with those of real data. The model was trained for 1500 epochs, achieving a reconstruction Fréchet Inception Distance (FID) score of 54.2.

Fig. 2 shows a qualitative comparison between real SVHN images and samples generated by the circular Kuramoto diffusion model. The generated samples exhibit recognizable digit-like structure and diversity across categories, indicating that the model learns meaningful representations while operating through a phase-based orientation field.

A. Phase Distribution Matching

To evaluate whether the model captures the underlying circular structure, we compare the distribution of phase values between real and generated data. The generated phase distribution follows the real distribution, indicating that the model learns the empirical statistics of the circular phase space. Minor deviations near the boundaries may arise from finite sampling and approximation errors in the wrapped diffusion process.

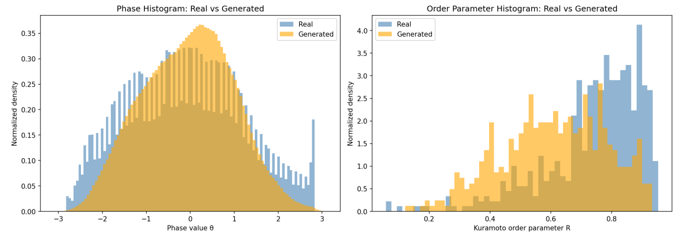

Fig. 3. Left: phase histogram comparison between real (blue) and generated (orange) samples. Right: distribution of the Kuramoto order parameter r for real (blue) and generated (orange) samples. The order parameter measures collective synchronization in the circular S^1 specialization of the vector Kuramoto framework.

B. Synchronization Statistics

To assess whether the model captures collective phase behavior, we analyze the distribution of the Kuramoto order parameter r, which is the circular analogue of the vector order parameter introduced in the hyperspherical formulation. Fig. 3 compares the distribution of r for real and generated samples. The generated distribution follows the real-data distribution, suggesting that the model captures not only local phase statistics but also global synchronization structure in the S^1 setting.

C. EEG Data Analysis

To evaluate the circular Kuramoto diffusion model on real-world neural phase data, we analyzed EEG recordings from the OpenNeuro dataset `ds004504`. In this application, each EEG channel phase is represented as a point on S^1, so the experiment tests the circular special case of the proposed hyperspherical framework.

Fig. 4 shows representative EEG traces from the validation set. Columns correspond to the three diagnostic groups (healthy controls, AD, FTD) and rows correspond to four scalp regions (frontal, temporal, parietal, occipital). In each panel the blue curve denotes the filtered EEG signal and the red dashed curve denotes the reconstruction from the phase-based diffusion model.

Fig. 5 compares the phase distributions of the real EEG signals and their reconstructions. Each panel corresponds to one diagnostic group, overlaying the empirical phase distribution of the real data with the phase distribution of the reconstructed signals. This comparison shows that the learned model preserves the circular structure of EEG phase values across different clinical populations.

IV. Conclusion

This work presented a Kuramoto-based diffusion framework for modeling oscillator-valued data on non-Euclidean state spaces. The model was formulated in a vector form on the hypersphere, where oscillator states are constrained to remain on S^{d-1} and synchronization corresponds to tangent-space alignment. The classical phase-based Kuramoto diffusion model on the circle arises as the S^1 special case, which is the setting used in the present experiments.

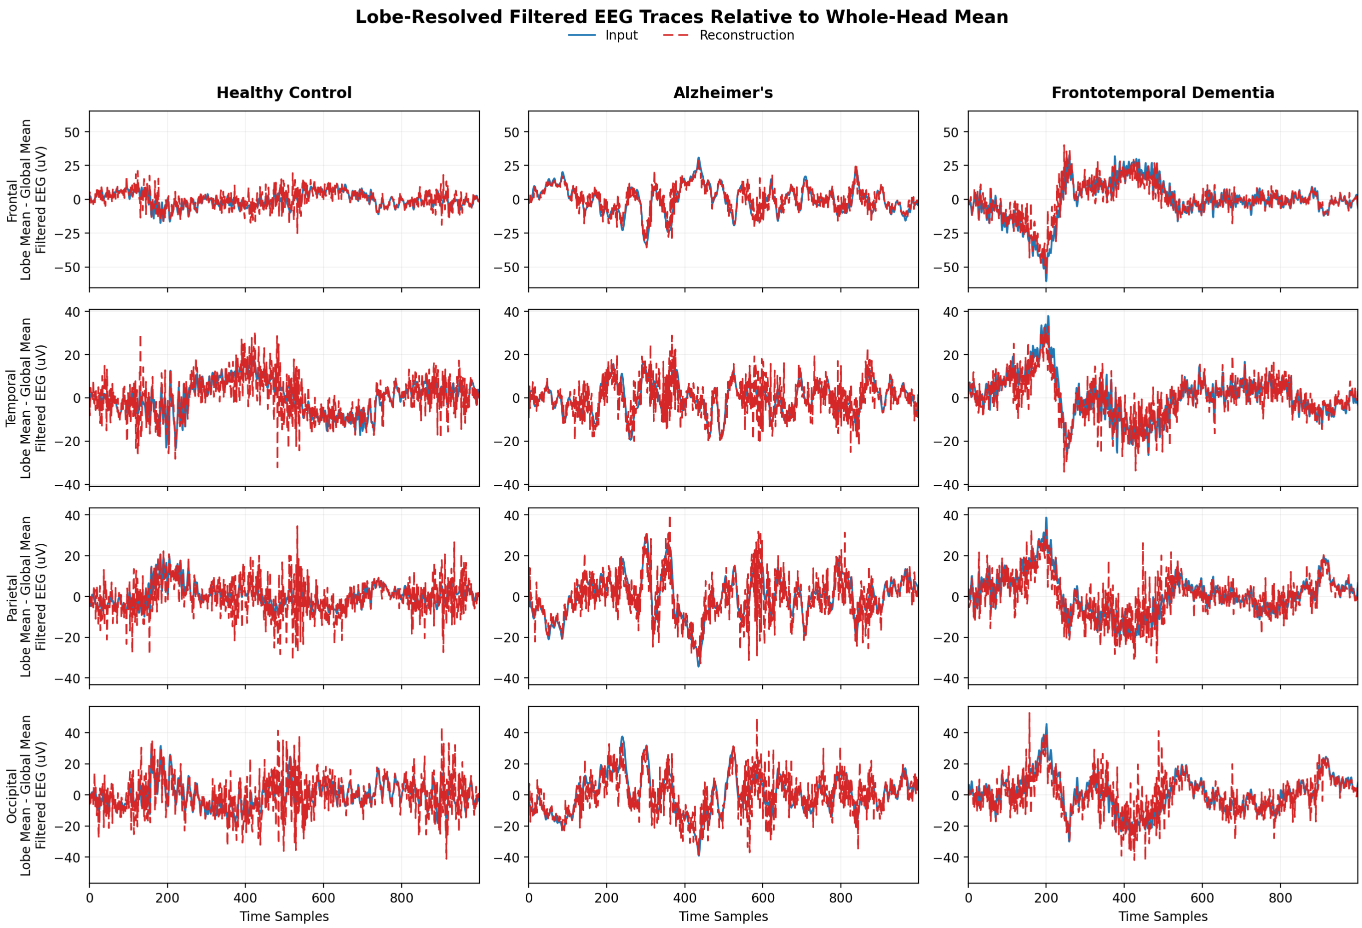

Fig. 4. Representative EEG traces for the three diagnostic groups. Columns correspond to healthy controls, Alzheimer's disease, and frontotemporal dementia; rows correspond to frontal, temporal, parietal, and occipital regions. In each panel the blue curve shows the filtered EEG signal and the red dashed curve shows the corresponding reconstruction from the phase-based diffusion model.

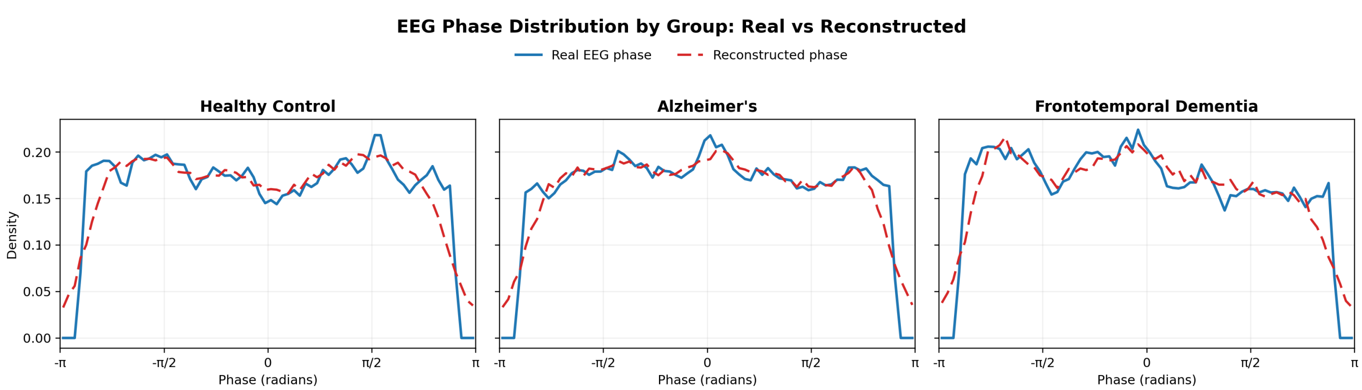

Fig. 5. Comparison of EEG phase distributions for real and reconstructed signals across diagnostic groups. Each panel corresponds to one group and overlays the empirical phase distribution of the real EEG data with that of the reconstructed signals.

In the circular setting, stochastic Kuramoto dynamics replace the standard Euclidean drift used in conventional diffusion models. This allows the forward and reverse processes to respect the periodic geometry of phase variables while incorporating synchronization-driven interactions. On the SVHN dataset, the model learns generative representations of image orientation fields and preserves both local phase statistics and global Kuramoto order-parameter structure. On EEG data, the model provides a generative description of neural phase dynamics across healthy controls, Alzheimer's disease patients, and frontotemporal dementia patients.

Several directions follow naturally from this formulation. Future work should empirically evaluate the full hyperspherical model for higher-dimensional oscillator states, including directional neural embeddings and other vector-valued phase representations. Incorporating heterogeneous coupling, anatomical or functional network structure, and spatial locality may improve realism for brain data, where interactions are not all-to-all. Learning data-driven coupling functions and diffusion schedules may further increase the flexibility and expressivity of Kuramoto-based diffusion models.

Acknowledgment

The author thanks Archbishop Mitty High School for providing the research environment in which this work was conducted.

References

[1] Y. Kuramoto, "Self-entrainment of a population of coupled nonlinear oscillators," in *Int. Symp. Mathematical Problems in Theoretical Physics*, Lecture Notes in Physics, vol. 39. Berlin: Springer, 1975, pp. 420–422.

[2] S. H. Strogatz, "From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators," *Physica D*, vol. 143, no. 1–4, pp. 1–20, 2000.

[3] M. Breakspear, "Dynamic models of large-scale brain activity," *Nature Neuroscience*, vol. 20, pp. 340–352, 2017.

[4] L. Yang *et al.*, "Diffusion models: A comprehensive survey of methods and applications," *ACM Computing Surveys*, vol. 56, no. 4, pp. 1–39, 2023.

[5] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, "Reading digits in natural images with unsupervised feature learning," in *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.

[6] Y. Song *et al.*, "Kuramoto orientation diffusion models," 2025, arXiv:2509.15328. [Online]. Available: `https://arxiv.org/abs/2509.15328`

[7] A. Miltiadous *et al.*, "A dataset of EEG recordings from: Alzheimer's disease, frontotemporal dementia and healthy subjects," OpenNeuro, 2024. [Online]. Available: `https://doi.org/10.18112/openneuro.ds004504.v1.0.8`

[8] W. Klimesch, "EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis," *Brain Research Reviews*, vol. 29, no. 2–3, pp. 169–195, 1999.

[9] G. Buzsáki, *Rhythms of the Brain*. Oxford: Oxford University Press, 2006.

[10] P. Fries, "A mechanism for cognitive dynamics: neuronal communication through neuronal coherence," *Trends in Cognitive Sciences*, vol. 9, no. 10, pp. 474–480, 2005.

[11] M. A. Lohe, "Quantum synchronization over quantum networks," *Journal of Physics A: Mathematical and Theoretical*, vol. 43, no. 46, p. 465301, 2010.

[12] S. Chandra, M. Girvan, and E. Ott, "Continuous versus discontinuous transitions in the D-dimensional generalized Kuramoto model: Odd D is different," *Physical Review X*, vol. 9, no. 1, p. 011002, 2019.

[13] K. V. Mardia and P. E. Jupp, *Directional Statistics*. Chichester, U.K.: John Wiley & Sons, 2000.

[14] E. P. Hsu, *Stochastic Analysis on Manifolds*, Graduate Studies in Mathematics, vol. 38. Providence, RI, USA: American Mathematical Society, 2002.