Kuramoto Orientation Diffusion Models

Yue Song[1], T. Anderson Keller[2], Sevan Brodjian[1], Takeru Miyato[3,4], Yisong Yue[1], Pietro Perona[1], Max Welling[4,5]

[1]Caltech [2]Harvard University [3]University of Tűbingen
[4]University of Amsterdam [5]CuspAI

NeurIPS 2025

Synchronization in Nature

bird flock

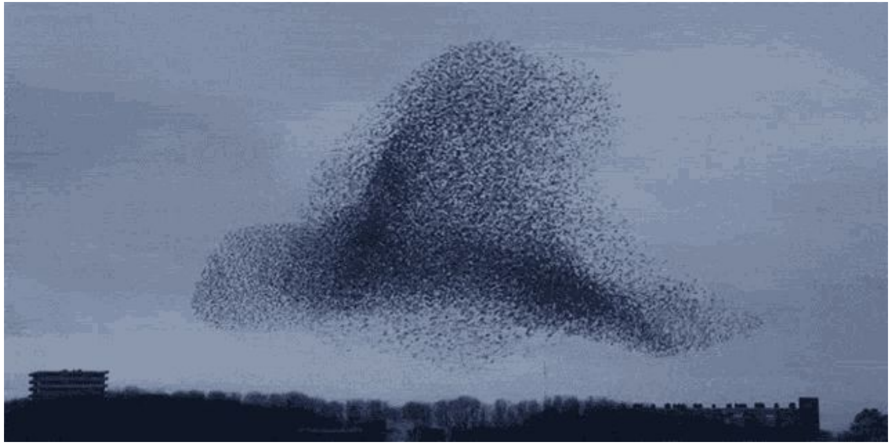

firefly

Phase synchronization in brains and physics:

Kuramoto model [1]: $$\frac{d\boldsymbol{\theta}_t^i}{dt} = \frac{1}{N}\sum_{j=1}^{N} K(t)\sin(\boldsymbol{\theta}_t^j - \boldsymbol{\theta}_t^i)$$

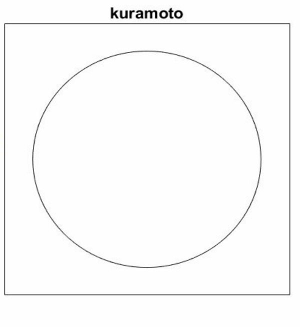

[1] Yoshiki Kuramoto. "Chemical turbulence". Springer, 1984.

Kuramoto Orientation Diffusion Models

Forward SDE:

$$\frac{d\boldsymbol{\theta}_t^i}{dt} = \frac{1}{N}\sum_{j=1}^{N} K(t)\sin(\boldsymbol{\theta}_t^j - \boldsymbol{\theta}_t^i) + K_{\text{ref}}(t)\sin(\psi_{\text{ref}} - \boldsymbol{\theta}_t^i) + \sqrt{2D_t}\boldsymbol{w}^i$$

phase coupling ***reference attraction***

$$\boldsymbol{\theta} = (\boldsymbol{\theta} + \pi) \ \texttt{mod} \ (2\pi) - \pi$$

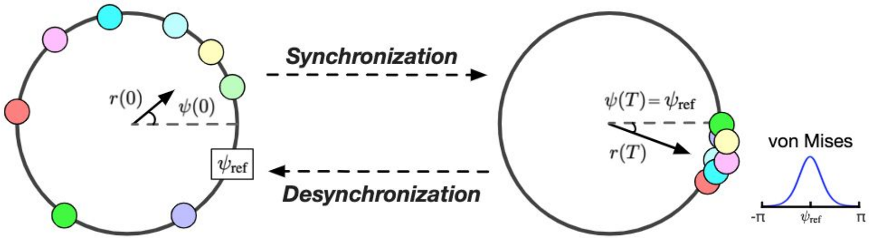

Terminal distribution (von Mises):

$$p_{\text{st}}(\theta) \approx \frac{1}{Z}\exp\left(\frac{K(T)r(T) + K_{\text{ref}}(T)}{D_T}\cos(\psi_{\text{ref}} - \theta)\right)$$

$r(T)$:average magnitude.

Learning the Score Function

Score-based generative models [1]:

$$\texttt{Forward-SDE:}\ d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x}, t)\, dt + g(t)\, d\boldsymbol{w}$$

$$\texttt{Reverse-SDE:}\ d\boldsymbol{x} = [\boldsymbol{f}(\boldsymbol{x}, t) - g^2(t) \boxed{\nabla_{\boldsymbol{x}} \log p(\boldsymbol{x})}]\, dt + g(t)\, d\bar{\boldsymbol{w}}$$

For linear drifts, $p(\boldsymbol{x}_t|\boldsymbol{x}_0)$ will be analytical Gaussian.

Kuramoto drift is highly non-linear; $p(\boldsymbol{x}_t|\boldsymbol{x}_0)$ is intractable.

Local Score Matching [2]. Approximation by MC sampling from the local transition kernel:

$$\nabla_{\boldsymbol{\theta}_t} \log p(\boldsymbol{\theta}_t) = \mathbb{E}_{\boldsymbol{\theta}_{t-1} \sim p(\boldsymbol{\theta}_{t-1}|\boldsymbol{\theta}_t)} \left[\nabla_{\boldsymbol{\theta}_t} \log p(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}) \right]$$

$$p(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}) = \mathcal{WN}\left(\boldsymbol{\theta}_{t-1} + \boldsymbol{f}(\boldsymbol{\theta}_{t-1}, t-1), 2D_{t-1}\boldsymbol{I}\right)$$

[1] Song, Yang, et al. "Score-based generative modeling through stochastic differential equations." ICLR (2021).
[2] Vincent, Pascal. "A connection between score matching and denoising autoencoders." Neural computation (2011).

Non-isotropic Noise Dynamics

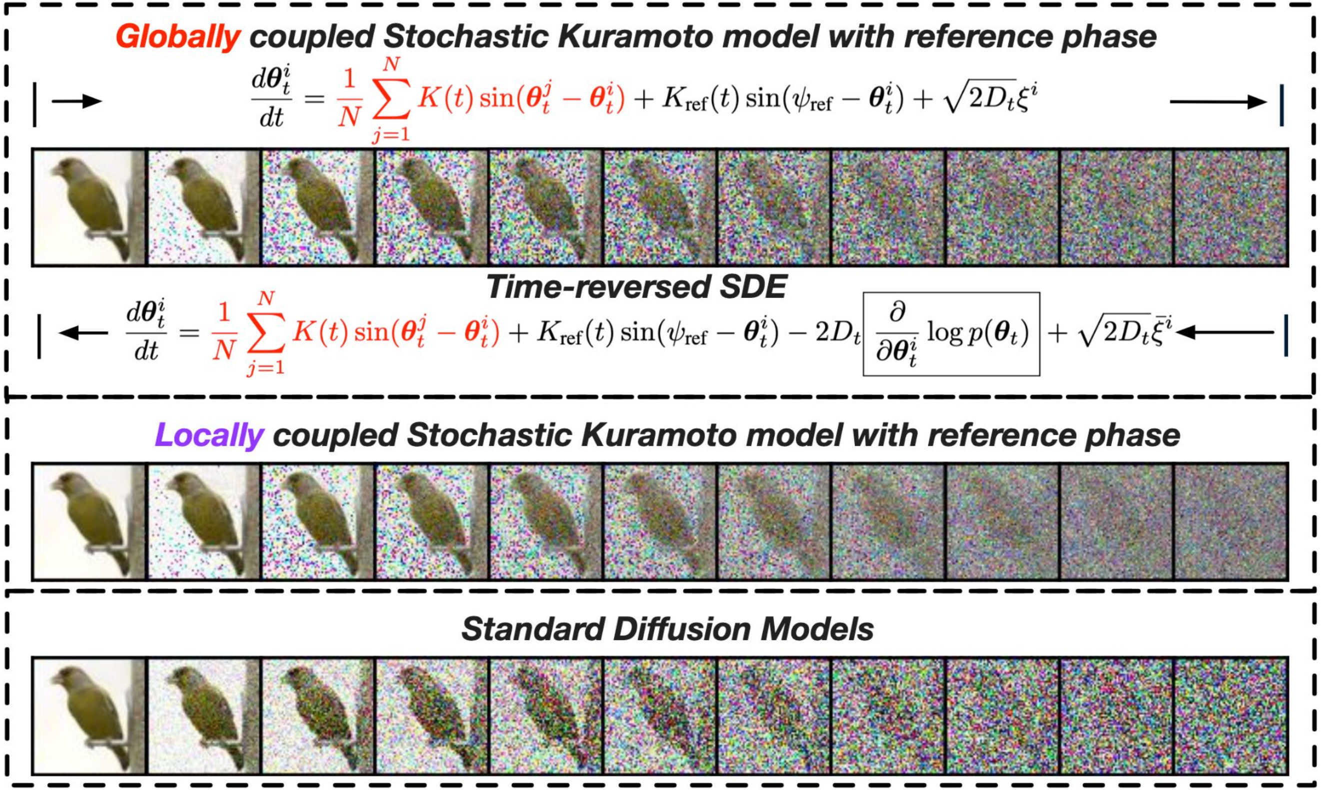

Learning to reverse Kuramoto model benefit more directional datasets!

Structured Destruction

SNR of the forward simulation

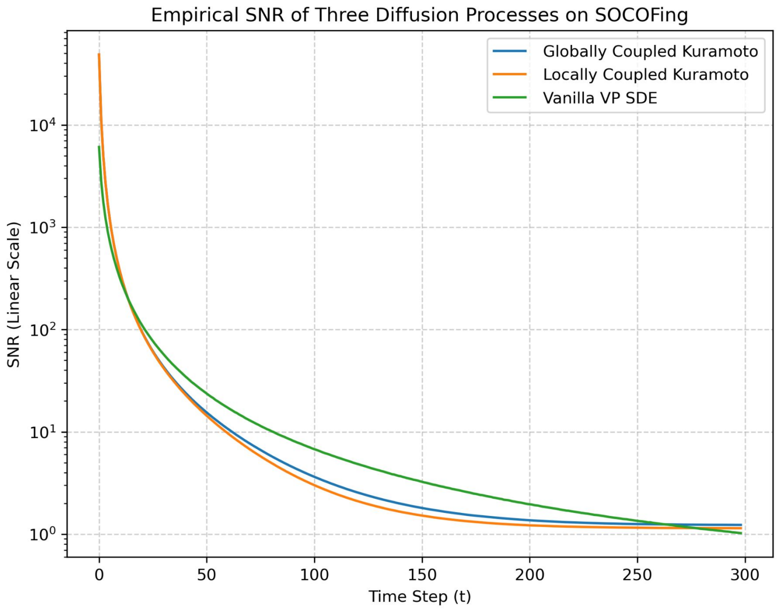

Synchronization dynamics:

Higher initial SNR: better structure preserving.

Faster cvgc: fewer steps for generation.

Forward and Backward Processes

FP (Diffusion Model [1])

FP (Ours)

Learned BP (Ours)

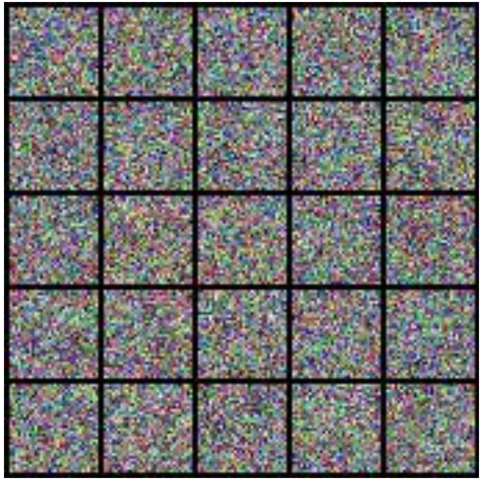

[1] Song, Yang, et al. "Score-based generative modeling through stochastic differential equations." ICLR (2021).

Orientation-dense Dataset

Brodatz Textures [1] (32x32)

Table 2: FID Scores (↓) on Brodatz texture dataset [1, 8].

Steps	100	300	1000
SGM [59]	38.33	22.40	20.37
Kuramoto Orientation Diffusion (Globally Coupled)	20.26	18.51	15.42
Kuramoto Orientation Diffusion (Locally Coupled)	18.47	15.93	14.19

100 steps

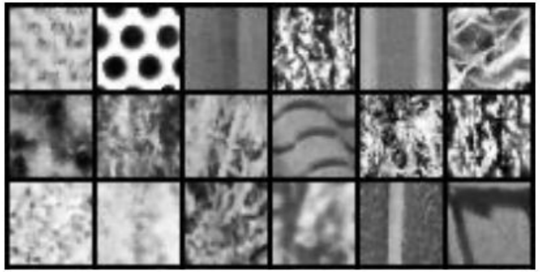

300 steps

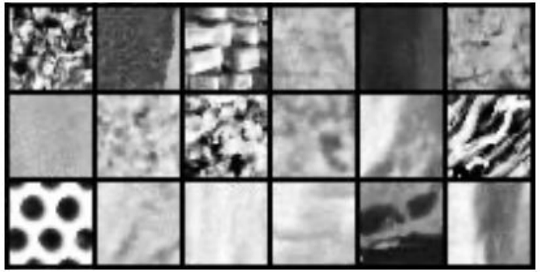

1000 steps

[1] Phil Brodatz. Textures: a photographic album for artists and designers. 1966.

Orientation-dense Dataset

SOCOFing Fingerprints [1] (96x96)

Table 1: FID Scores (↓) on SOCOFing fingerprint dataset [55].

Diffusion Steps	100	300	1000
SGM [59]	104.92	62.66	23.84
Kuramoto Orientation Diffusion (Globally Coupled)	74.41	47.93	20.64
Kuramoto Orientation Diffusion (Locally Coupled)	67.49	43.57	18.75

100 steps

300 steps

1000 steps

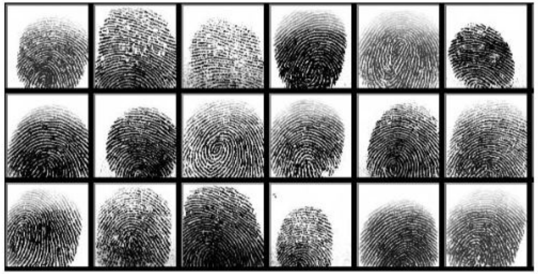

[1] Yahaya Isah Shehu et al. Sokoto coventry fingerprint dataset. arXiv (2018).

Orientation-dense Dataset

Ground Terrain [1] (128x128)

Table 3: FID Scores (↓) on the ground terrain dataset [67].

Diffusion Steps	100	300	1000
SGM [59]	114.90	56.72	33.79
Kuramoto Orientation Diffusion (Globally Coupled)	101.65	54.17	33.56
Kuramoto Orientation Diffusion (Locally Coupled)	92.86	49.68	30.62

100 steps

300 steps

1000 steps

[1] Jia Xue, Hang Zhang, and Kristin Dana. Deep texture manifold for ground terrain recognition. CVPR. (2018).

General Image Dataset

CIFAR10 [1] (32x32)

Table 4: FID Scores (↓) on CIFAR10 [34].

Diffusion Steps	100	300	1000
SGM [59]	38.04	25.76	3.17
Kuramoto Orientation Diffusion (Globally Coupled)	29.96	25.83	11.58
Kuramoto Orientation Diffusion (Locally Coupled)	28.17	24.86	10.79

100 steps

300 steps

1000 steps

[1] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. (2009).

Earth and Climate Sciece Dataset (2D Sphere)

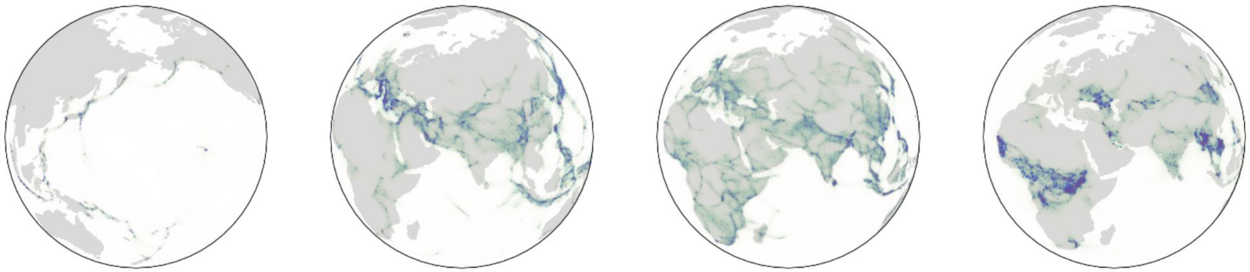

(a) Volcano (b) Earthquake (c) Flood (d) Fire

Table B: Test NLL on Earth and climate science datasets averaged across 5 runs.

Dataset	Volcano	Earthquake	Flood	Fire
Riemannian CNF [10]	-6.05±0.61	0.14±0.23	1.11±0.19	-0.80±0.54
Moser Flow [13]	-4.21±0.17	-0.16±0.06	0.57±0.10	-1.28±0.05
CNF Matching [1]	-2.38±0.17	-0.38±0.01	0.25±0.02	-1.40±0.02
Riemannian score-based [4]	-4.92±0.25	-0.19±0.07	0.48±0.17	-1.33±0.06
Riemannian diffusion model [8]	-6.61±0.96	-0.40±0.05	0.43±0.07	-1.38±0.05
Riemannian flow matching [3]	-7.93±1.67	-0.28±0.08	0.42±0.05	-1.86±0.11
Our Kuramoto orientation diffusion model	-5.18±0.17	-0.18±0.06	0.49±0.18	-1.44±0.05

Complex diffusion models

Shallow water Eq. of wind velocity

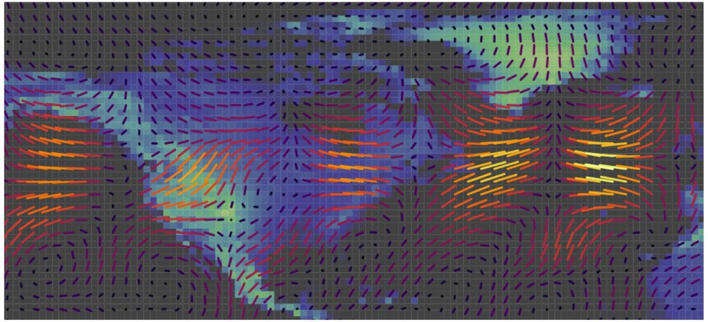

Navier-Stokes Eq. of fluid vorticity

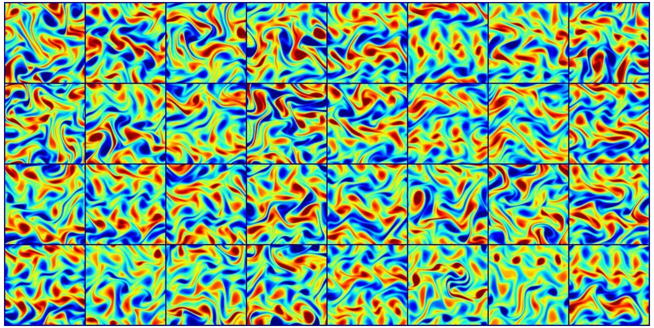

Data is always represented in Cartesian coordinates $(\boldsymbol{v}_x, \boldsymbol{v}_y)$.

What if we use polar form $\boldsymbol{s}_t = \boldsymbol{r}_t e^{i\boldsymbol{\theta}_t}$?

Magnitude: Standard Diffusion **Phase**: Kuramoto Diffusion

Preliminary Spectral Evaluation

Navier-Stokes Fluids

Diffusion Model

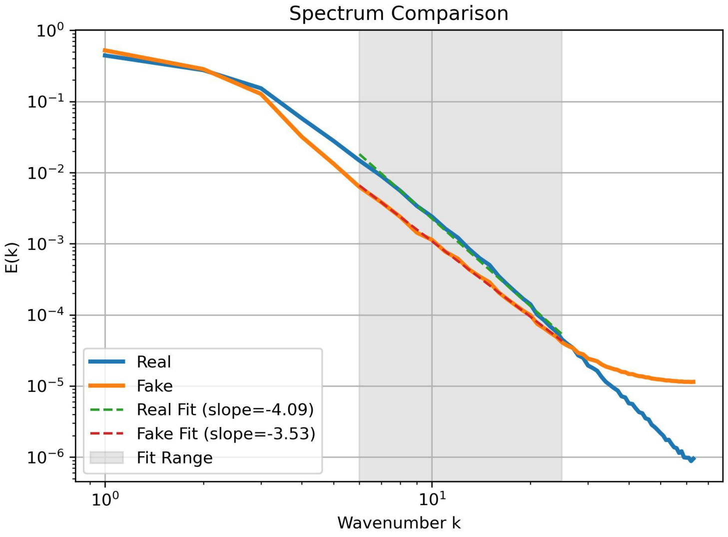

Complex Diffusion Model

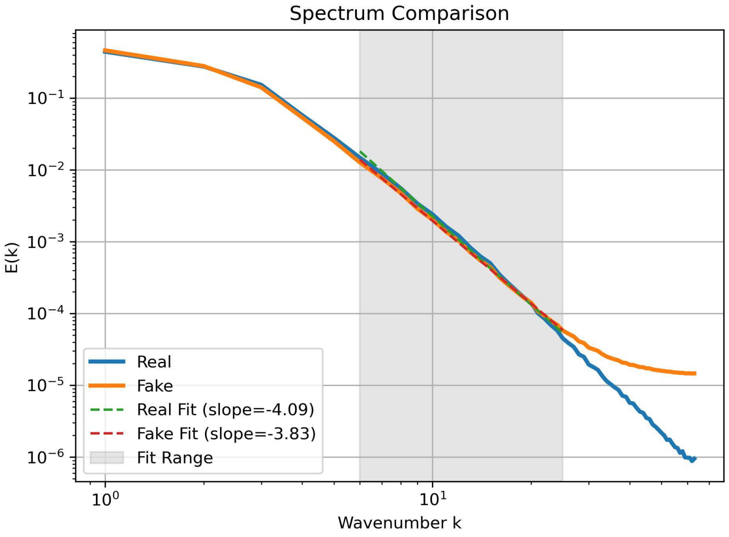

Kuramoto Orientation Diffusion Models

Synchronization inductive bias:

- Non-isotropic noise dynamics;
- Structured destruction;
- Hierachical generation process.

Empirical results:

- Higher sampling efficiency;
- Large gain on orientation-dense datasets;
- Competitive results on general benchmarks;
- Potential for vector-valued generative modelling.