

Hybrid generative models in neuroscience

Whole brain dynamics

Daniel López-Montero

August 25, 2026

Chair for Dynamics, Control, Machine Learning and Numerics, FAU Erlangen-Nürnberg

joint work with L. Liverani, E. Zuazua and X. Kobeleva

Why model a brain?

A century of brain models

Brain models: from physics to data-driven

The fitting problem

Results

Why model a brain?

Why model a brain?

Concrete uses of a brain model:


- **Understand.** Uncover the brain's anatomy and physiology.
- **Diagnose.** Detect and identify neurological illnesses early.
- **Treat.** Test a stimulation protocol *in silico* first.

Requirements:

1. **Data fidelity.** Produce *descriptors* similar to those of the real brain.
2. **Interpretability.** Increase trustworthiness and adoption among neuroscientists.
3. **Personalisation.** One model per brain, trainable from the few recordings available for a single subject.

A century of brain models

A century of brain models



fMRI in one picture



The data, formally

Human Connectome Project

$$\mathcal{D} = \{x^{(k)}\}_{k=1}^N, \quad x^{(k)} : [0, T] \rightarrow \mathbb{R}^{100},$$


$T_{\text{obs}} = 1200$ samples, $\Delta t = 0.72$ s.

Preprocessing

1. Standardise each region $(x - \mu)/\sigma$;
2. filter to $[0.008, 0.08]$ Hz;
3. apply the Hilbert transform $\Rightarrow$ analytic signal.

Brain models: from physics to data-driven


Different scales of brain models



Two families of brain models


Neural mass models (ODE/SDE) — finitely many states in $\mathbb{R}^n$ or $\mathbb{C}^n$.

$$\dot{z}_i = f(z_i) + \sum_j C_{ij} g(z_i, z_j)$$




Neural field models (PDE) — differential equations on a cortical sheet $\Omega \subset \mathbb{R}^3$:

$$\partial_t \psi = -\psi + \int_{\Omega} \omega(x - x') f(\psi(t, x')) dx'$$



Stuart–Landau oscillator

$$\dot{z} = (a + i\omega - |z|^2)z \quad \Longleftrightarrow \quad \begin{cases} \dot{r} = r(a - r^2) \\ \dot{\phi} = \omega \end{cases}$$




The Coupled Hopf model

$$dz_i = \underbrace{\left[(a + i\omega_i - |z_i|^2) z_i \right]}_{\text{local oscillator}} + \underbrace{G \sum_{j=1}^n C_{ij} (z_j - z_i)}_{\text{coupling}} dt + \sigma dW_i$$

where

- a — bifurcation parameter (excitability);
- ω_i — intrinsic frequency;
- G — one global coupling gain;
- C — *fixed* anatomy from diffusion MRI;
- σ — noise amplitude;
- W_i — independent complex Brownian motions.



Structural connectome C from MRI

The data-driven extreme: Neural SDEs

Neural Stochastic Differential Equations (SDEs)

$$dz(t) = f_{\theta}(z(t)) dt + g_{\phi}(z(t)) dW(t), \quad f_{\theta}, g_{\phi} : \mathbb{C}^n \rightarrow \mathbb{C}^n \text{ neural networks}$$

Fitting. Backpropagate through the Euler–Maruyama solver to get $\nabla_{\theta} \mathcal{L}$.

$$\theta_{k+1} = \theta_k - \eta \nabla_{\theta} \mathcal{L}(\hat{z}_{\theta}, z)$$

The cost per update is independent of $\dim \theta$.

- ✓ universal approximation of drift *and* diffusion;
- ✓ computationally efficient;
- ✗ no oscillatory prior $\Rightarrow$ data hungry;
- ✗ no parameter has a meaning.

So gradients solve the scaling problem. But against what objective?

The fitting problem

A 50-year-old habit and why it fails

The default objective in trajectory fitting is the L^2 misfit

$$\mathcal{L}(\theta) = \mathbb{E} \|\hat{z}_\theta - z\|_{L^2(0, T)}^2.$$

A 50-year-old habit and why it fails

The default objective in trajectory fitting is the L^2 misfit

$$\mathcal{L}(\theta) = \mathbb{E} \|\hat{z}_\theta - z\|_{L^2(0,T)}^2.$$

Observation

Let z be the empirical process and $\hat{z}_\theta$ the model, driven by a Brownian motion. Then

$$\mathbb{E} \|\hat{z}_\theta - z\|^2 = \mathbb{E} \|\hat{z}_\theta\|^2 + \mathbb{E} \|z\|^2 - 2 \mathbb{E} \langle \hat{z}_\theta, z \rangle = \mathbb{E} \|\hat{z}_\theta\|^2 + \mathbb{E} \|z\|^2,$$

since $\hat{z}_\theta$ is independent of z and the signal is standardised ($\mathbb{E} z = 0$).

A 50-year-old habit and why it fails

The default objective in trajectory fitting is the L^2 misfit

$$\mathcal{L}(\theta) = \mathbb{E} \|\hat{z}_\theta - z\|_{L^2(0,T)}^2.$$

Observation

Let z be the empirical process and $\hat{z}_\theta$ the model, driven by a Brownian motion. Then

$$\mathbb{E} \|\hat{z}_\theta - z\|^2 = \mathbb{E} \|\hat{z}_\theta\|^2 + \mathbb{E} \|z\|^2 - 2 \mathbb{E} \langle \hat{z}_\theta, z \rangle = \mathbb{E} \|\hat{z}_\theta\|^2 + \mathbb{E} \|z\|^2,$$

since $\hat{z}_\theta$ is independent of z and the signal is standardised ($\mathbb{E}z = 0$).

Consequence

If $\theta^* := \arg \min_{\theta} \mathcal{L}(\theta)$, then $\hat{z}_{\theta^*} \approx 0$: the L^2 -optimal generative model is degenerate.

The right objects: statistics of the law, not paths

Replace path matching by matching a family of **functionals of the law** $\Phi_1, \dots, \Phi_m$:

$$\mathcal{L}(\theta) = \sum_{k=1}^m w_k d_k\left(\Phi_k(\hat{z}_\theta), \Phi_k(z)\right) + \lambda \mathcal{R}(\theta).$$

We ask two things of every Φ_k :

- **Noise-invariant** — Φ_k is a property of the *law* of z , not of the path drawn, so the objective is well posed for a stochastic generative model.
- **Differentiable** — Φ_k is a differentiable function, so we can backpropagate through it.

Network structure

- FC_{ij} — region correlation matrix
- FC_{ij}^ϕ — phase synchrony (phFC)

Node-level features

- mean amplitude, mean frequency, metastability

Network dynamics

- $\text{FCD}(t, t')$ — temporal correlation
- $\text{phFCD}(t, t')$ — temporal phase synchrony

Regulariser

- Finite-dimensional matching: two-time joint distributions $(\hat{z}_\theta(t_1), \hat{z}_\theta(t_2))$


Example: Functional Connectivity (FC)

Standardise each region on $[0, T]$:

$$\tilde{x}_i := \frac{x_i - \mu_i}{\sigma_i}, \quad \text{with} \quad \mu_i := \frac{1}{T} \int_0^T x_i \, dt, \quad \sigma_i^2 := \frac{1}{T} \int_0^T (x_i - \mu_i)^2 \, dt$$

then calculate the correlation between every pair of regions i, j :

$$\text{FC}_{ij}(z) := \frac{1}{T} \int_0^T \tilde{x}_i(t) \tilde{x}_j(t) \, dt.$$



Physical vs. data-driven models


	Coupled Hopf	Neural SDE
Data fidelity	✗ too rigid	✓ flexible
Data efficiency	✓ few parameters	✗ data hungry
Interpretability	✓ physical model	✗ black box

These two models sit at opposite extremes of the same spectrum.

Can we build a model that keeps the best of both worlds?


Three architectures on one axis

a Coupled Hopf



$$\dot{z}_i = \begin{bmatrix} (a + i\omega_i - |z_i|^2)z_i \\ +G \sum_{j=1}^n C_{ij}(z_j - z_i) \end{bmatrix}$$


b Hybrid Hopf



$$m_i = \sum_{j=1}^n C_{ij}(k_i - k_j), \quad k_j = \phi_{\text{key}}(z_j)$$

$$\dot{z}_i = [(a + i\omega_i - |z_i|^2)z_i + \psi(m_i, z_i)]$$

c Neural SDE



$$dz = f_{\theta}(z)dt + g_{\phi}(z)dW_t$$

Physics-based

More interpretable

Data-driven

More flexible

Results

Model comparison

Model	network structure			network dynamics		
	FC corr $\uparrow$	phFC corr $\uparrow$	FC ARI $\uparrow$	FCD KS $\downarrow$	phFCD KS $\downarrow$	Meta $ \Delta $ $\downarrow$
Coupled Hopf	0.911	0.783	0.124	0.235	0.254	9.5×10^{-3}
Neural SDE	0.915	0.911	0.328	0.049	0.305	6.7×10^{-3}
Hybrid Hopf	0.916	0.915	0.210	0.082	0.052	1.3×10^{-2}

Summary of the trade-off — and where it moved

	Coupled Hopf	Hybrid Hopf	Neural SDE
Data fidelity, static	~	✓	✓
Data fidelity, dynamic	✗	✓	~
Interpretable parameters	✓	✓	✗
Uses the connectome	✓	✓	✗
Data efficiency ($N = 10$)	✓	✓	✗
Personalisation ($N = 1$)	✗	✓	✓
Free parameters	2 – 5	~ 2 200	~ 220 000

Multi-objective optimisation

$\mathcal{L} = \sum_k w_k d_k$ with fixed w_k picks one point on a Pareto front.

Control

The clinical goal is stimulation: add $Bu(t)$ to the drift and ask for controllability, optimal protocols, and closed-loop feedback on a learned stochastic system. Does the learned model preserve the *sensitivity* $\partial_u y$ of the real brain?

Mean-field model and asymptotic analysis

As the number of neurons goes to infinity, $n \rightarrow \infty$, mass models reconnect to the neural field equations they came from. What properties does the system have when $t \rightarrow \infty$?

Thank you

`github.com/DCN-FAU-AvH/neuroscience-modeling`