MAGNET: Manifold-Aware Graph Diffusion Network for Connectome Generation

Protyay Dey[1*] Ayush Roy[1*] Hyuna Cho[2] Won Hwa Kim[3] Vishnu Suresh Lokhande[1]

[1] University at Buffalo [2] NIH [3] POSTECH

* Equal contribution: Protyay Dey and Ayush Roy

NEURAL INFORMATION PROCESSING SYSTEMS NeurIPS 2026

PROJECT WEBSITE

THE PROBLEM

Functional connectomes encode interactions among brain regions. Valid correlation matrices alone do not ensure **valid correlation matrices** and **realistic clinical structure.**

Coarse diagnostic labels conceal **disease stage, subtype**, and **patient-to-patient variation.** Generators can **average over subgroups** and lose **minority patterns.**

Our goal: generate **valid, diverse connectomes** that preserve **within-class clinical heterogeneity**, even when training labels are only coarse.

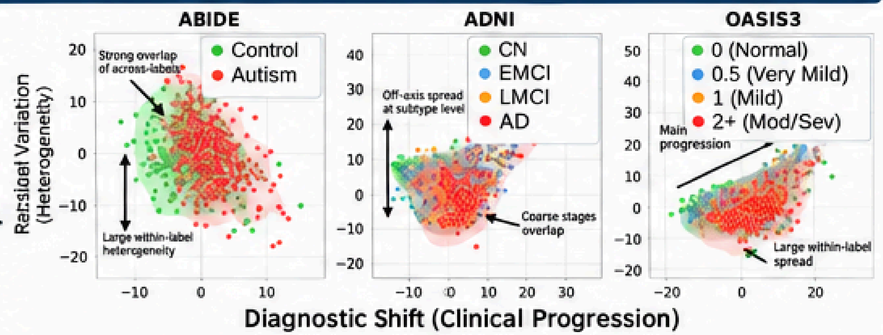

MAGNET FRAMEWORK

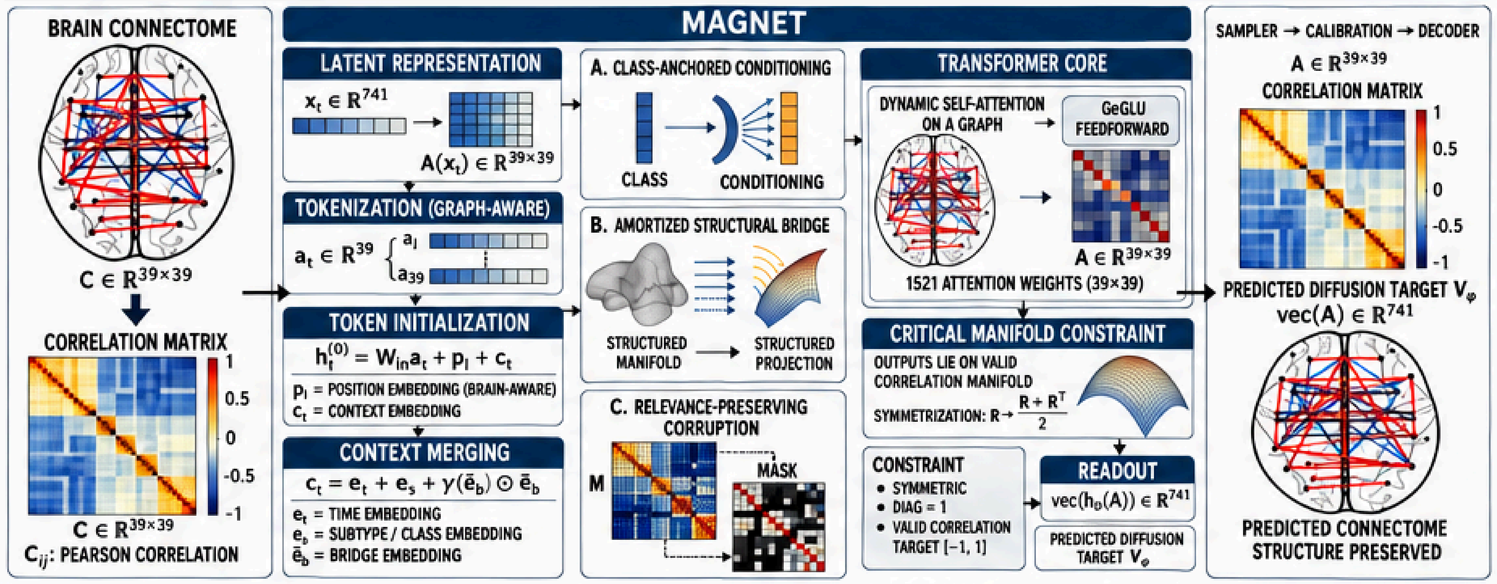

CORE COMPONENTS

A. CLASS-ANCHORED CONDITIONING

Model latent subtypes within each coarse clinical class; BIC selects mixture size.

$$p(x \mid y) \approx \sum_s \pi_{y,s} p(x \mid y, s)$$

Subtype-aware denoising prevents collapse toward a coarse class mean.

B. AMORTIZED STRUCTURAL BRIDGE

Predict a 16-dimensional structural target directly from each noisy latent state.

$$b_t = B_\psi(x_t, t, s)$$

$$\mathcal{L}_{\text{bridge}} = \|B_\psi(x_t, t, s) - b^*(x_0, s, t)\|_2^2$$

No clean-state inference at each step.

C. RELEVANCE-PRESERVING CORRUPTION

Protect condition-relevant coordinates by reducing early diffusion noise.

$$m_{t,j} = |d_{t,j}| / (\max_k |d_{t,k}| + \varepsilon)$$

$$\sigma_{t,j}^2 = \text{clip}(1 - \kappa\rho(t) m_{t,j}, \sigma_{\min}^2, 1)$$

Per-coordinate noise scheduling.

CLASS-CONDITIONAL GENERATION

Cohort	Method	α,β-F1 ↑	Time (s) ↓
ABIDE	MAGNET	**0.68**	0.38
	DiffeoCFM	0.62	9.04
	DiffeoGauss	0.32	**0.06**
	TriangCFM	0.20	7.16
	GDSS-proj	0.00	1.05
ADNI	MAGNET	**0.76**	**0.86**
	DiffeoCFM	0.71	36.08
	DiffeoGauss	0.04	1.42
	TriangCFM	0.23	21.70
	GDSS-proj	0.45	1.70
OASIS-3	MAGNET	**0.57**	**0.63**
	DiffeoCFM	0.47	23.24
	DiffeoGauss	0.39	1.88
	TriangCFM	0.23	15.29
	GDSS-proj	0.01	1.44

DOWNSTREAM CAS

ABIDE	
CAS F1 ↑	0.56
ROC-AUC ↑	0.64
ADNI	
CAS F1 ↑	0.61
ROC-AUC ↑	0.57
OASIS-3	
CAS F1 ↑	0.56
ROC-AUC ↑	0.68

Classifier-based utility

Takeaway: MAGNET achieves the best overall fidelity–coverage trade-off across the clinical cohorts while remaining fast to sample.

ZERO-SHOT SUBCLASS RECOVERY

Cohort	Method	α,β-F1 ↑	λ-MSE ↓
ADNI	MAGNET	**0.35**	**6.08**
	DiffeoCFM	0.33	14.30
	DiffeoGauss	0.01	11.37
	TriangCFM	0.25	9.16
	GDSS-proj	0.20	13.47
OASIS-3	MAGNET	**0.28**	**1.61**
	DiffeoCFM	0.26	1.93
	DiffeoGauss	0.21	2.24
	TriangCFM	0.26	1.75
	GDSS-proj	0.00	4.00

CAS + CLASS BALANCE

ADNI	
CAS macro-F1 ↑	0.28
Class DKL ↓	0.05
Class DB ↓	0.01
OASIS-3	
CAS macro-F1 ↑	0.36
Class DKL ↓	0.02
Class DB ↓	0.01

DKL: class KL divergence
DB: Bhattacharyya distance

Takeaway: Even when trained only with coarse labels, MAGNET better recovers hidden clinical subclasses and preserves clinically meaningful structure.

CORTICAL HUB ALIGNMENT

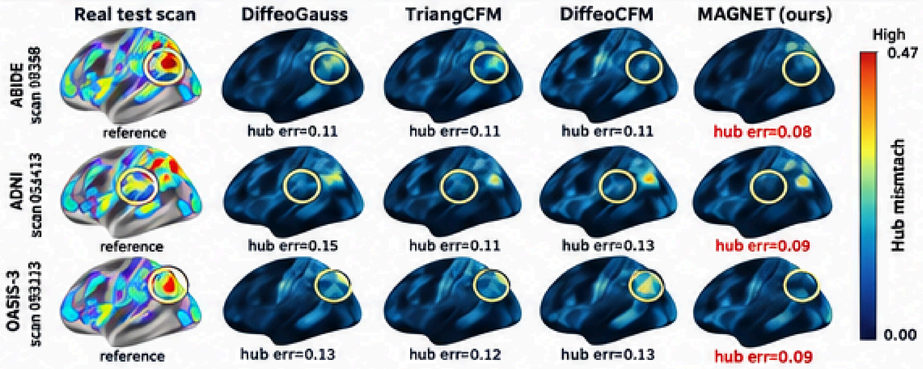

MAGNET achieves the lowest cortical hub mismatch across all three cohorts.