CONTINUOUS NEURAL FIELDS FOR FUNCTIONAL CONNECTIVITY

Anonymous authors
Paper under double-blind review

ABSTRACT

Brain Functional connectivity (FC) is typically represented after discretizing the brain into a predefined set of regions, making the resulting representation dependent on the chosen parcellation and spatial scale. We introduce a continuous representation of resting-state functional organization in which each cortical location is associated with a low-dimensional functional embedding, and connectivity is induced by similarity between embeddings. We parameterize this representation with implicit neural representations (INRs), first learning a population functional field that captures shared organization and then adapting it to each individual through constrained subject-specific deviations. The resulting continuous fields accurately represent functional connectivity while capturing reproducible individual variation across independent acquisitions. Beyond reconstruction, the differentiable representation allows cortical regions to serve as spatial readouts of the field, whose location and geometry can be optimized for downstream analyses. Across three cohorts, optimized regions yield stronger held-out associations with age and sex than the best connections of coarse and fine parcellations in every setting, and the advantage extends to behavioral and cognitive phenotypes. Together, these results establish continuous functional fields as an alternative to parcellation-defined FC representations, supporting individualized modeling and spatially flexible analysis within a common framework.

1 INTRODUCTION

Resting-state functional MRI (rs-fMRI) provides a window into the intrinsic functional organization of the human brain. Statistical dependencies between spontaneous blood-oxygen-level-dependent (BOLD) fluctuations, commonly summarized as functional connectivity (FC), have revealed reproducible large-scale functional systems and substantial variation in their organization across individuals (Biswal et al., 1995; Fox et al., 2005; Yeo et al., 2011; Finn et al., 2015). FC has consequently become a standard representation for studying individual differences, behavior, development, and neurological and psychiatric conditions (Smith et al., 2013; Shen et al., 2017).

In practice, FC is typically represented only after discretizing the brain into predefined regions based on a functional atlas (Figure1, left). However, many functional atlases are available, ranging from a few large-scale functional networks to fine-grained parcellations comprising thousands of regions (Yeo et al., 2011; Schaefer et al., 2018), but there is no consensus on which atlas or spatial scale is most appropriate for a given study. Once an atlas is chosen, individualized parcellations adapt these boundaries across subjects (Salehi et al., 2018), while supervised approaches optimize them for downstream endpoints (Michel et al., 2012; Xu et al., 2025; Hannum & Lopez, 2026). Nevertheless, the resulting representation of FC remains a finite, discrete region-by-region matrix. This raises a more fundamental question: *does FC need to be discretized in the first place?*

Indeed, the biological brain is not a collection of discrete parcels. Functional organization varies continuously across cortical space, while individual differences in brain function may emerge at diverse spatial scales and extend across atlas boundaries. We posit that this mismatch between *continuous* brain organization and *discrete* brain parcellation represents a fundamental bottleneck for interrogating the brain's functional organization.

Implicit neural representations (INRs) provide a natural framework for replacing sampled signals with continuous functions of their coordinates (Park et al., 2019; Sitzmann et al., 2020; Tancik et al.,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

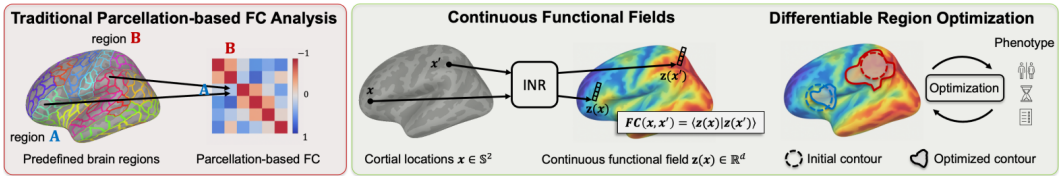


Figure 1: Unlike conventional parcellation-based FC, we represent functional organization as a continuous field over the cortex, enabling connectivity to be queried at arbitrary locations. This formulation enables differentiable spatial readouts that can be optimized for downstream variables.

2020). Rather than storing values at fixed spatial locations, coordinate networks parameterize an underlying signal continuously and can be evaluated independently of its original sampling grid. This paradigm has recently been introduced to functional neuroimaging by Chen et al. (2026), who model stimulus-evoked fMRI responses as a continuous function of anatomical coordinates. Their Neural Response Function predicts the BOLD response to a visual stimulus at arbitrary locations in standardized brain space, moving neural encoding beyond fixed voxel grids. In resting-state fMRI, however, there is no external stimulus whose local response can be predicted; the object of interest is instead the intrinsic functional relationship among spontaneous activity across brain regions.

Here, we propose a different paradigm for functional connectivity analysis, in which INRs capture the resting-state functional organization as a continuous *functional field* over the cortical surface (Figure 1, right). Each location is associated with a low-dimensional functional embedding whose geometry induces connectivity between arbitrary locations. We capture shared and individual functional organization through a population field and subject-specific adaptations, yielding individualized continuous representations without committing to a predefined parcellation.

Importantly, this continuous representation fundamentally changes how cortical regions can be used in downstream brain–phenotype analyses, where conventional approaches test associations between a phenotype and FC defined over predefined regions. In our framework, however, regions can be placed at arbitrary cortical locations to aggregate the underlying field into region-level functional representations. Because this operation is *differentiable* with respect to the regions' locations and shapes, the regions themselves can be *optimized* for downstream phenotype-association analyses. Our contributions are:

1. **Continuous functional connectivity.** We represent individual resting-state functional organization as a continuous low-dimensional neural field, replacing atlas-specific FC matrices with a representation that can be queried across spatial locations and scales.
2. **Population-to-individual modeling.** We capture subject-specific functional organization as continuous adaptations of a shared population field, providing a common representation while preserving individual variation.
3. **Differentiable spatial discovery.** We use the continuous field to define and optimize the locations and shapes of cortical regions whose connectivity shows maximal associations with phenotypic variables, without being constrained by predefined atlas boundaries.

2 Related Work

Functional connectivity and spatial parcellation. The most common way to analyze resting-state fMRI is to study FC among brain regions divided according to an atlas defined on the population level (Biswal et al., 1995; Yeo et al., 2011; Schaefer et al., 2018). Beyond population parcellation, individualized methods adapt functional boundaries to subject-specific organization (Salehi et al., 2018). These parcellations can even adapt to downstream analyses. For example, Michel et al. (2012) use supervised prediction to select spatially contiguous voxel clusters, Xu et al. (2025) optimize a discrete node partition according to functional homogeneity and associations between inter-node FC and behavioral phenotypes, and Hannum & Lopez (2026) learn task-specific parcellations for connectivity-based classification. Our objective is fundamentally different: we learn a single, endpoint-independent continuous functional representation, that serves as a common founda-

tion for downstream phenotype-association analyses, allowing spatial readouts to adapt to different questions without being constrained by a predefined discretization.

Individualized functional organization. Functional organization exhibits substantial and reproducible inter-individual variability (Mueller et al., 2013; Finn et al., 2015; Gordon et al., 2017), motivating individualized parcellations and subject-specific network representations (Salehi et al., 2018). Individual connectivity patterns are sufficiently stable to identify subjects across scanning sessions and are associated with behavioral phenotypes (Finn et al., 2015). Our approach similarly separates shared from individual functional organization, but represents both continuously.

Implicit neural representations and continuous brain modeling. INRs parameterize signals as neural functions of their coordinates and have been widely used to represent continuous images, geometry, and other spatial signals (Park et al., 2019; Sitzmann et al., 2020; Tancik et al., 2020). Chen et al. (2026) recently introduced this perspective to fMRI by modeling stimulus-evoked BOLD responses continuously over anatomical coordinates, enabling spatially aware neural encoding beyond fixed voxel grids. We instead consider resting-state fMRI, where the target is not a local response to an external stimulus but the intrinsic relationship between spontaneous activity across locations. We therefore use a continuous embedding field to represent functional organization, from which connectivity and spatially flexible region-level readouts can be derived.

3 METHODS

Our framework represents resting-state functional organization as a continuous field over the cortical surface (Figure 1). We first derive low-dimensional functional embeddings from rs-fMRI BOLD signals sampled at discrete cortical vertices for each individual. These embeddings are then decomposed to identify a population-level embedding that captures shared functional structure. The population embedding is parameterized continuously using implicit neural representations, yielding a continuous population field, while subject-specific residual fields are fit for each individual to capture individualized deviations from the population structure. Finally, the continuous representation enables evaluation of functional connectivity across differentiable spatial regions, whose locations and geometries can be optimized for downstream analyses.

3.1 CONTINUOUS MODELING OF FUNCTIONAL CONNECTIVITY

Functional connectivity as a continuous field. We seek a continuous representation of functional organization over the cortical surface, without first partitioning the cortex into a finite set of regions. Let $\mathbb{S}^2$ denote the cortical surface. We propose to represent each cortical location $x \in \mathbb{S}^2$ by a low-dimensional functional embedding $z(x) \in \mathbb{R}^d$, such that connectivity between two locations is induced from their inner product,

$$\mathrm{FC}(x, x') = \langle z(x), z(x') \rangle . \tag{1}$$

A single embedding field $z(x)$, therefore, induces the entire connectivity structure continuously over cortical space and can be evaluated at arbitrary locations. We parameterize this field with an implicit neural representation (INR): a coordinate-based network that continuously maps cortical position to its functional embedding. Writing $\mathrm{enc}(\cdot)$ for a Fourier-feature encoding of position, the INR outputs, for $x \in \mathbb{S}^2$:

$$z(x) = \mathrm{MLP}(\mathrm{enc}(x)) \in \mathbb{R}^d. \tag{2}$$

Learning the continuous field from sampled fMRI. Next, we introduce an efficient way to learn the continuous functional field from the rs-fMRI BOLD signal measured at a finite set of cortical locations. Let $\tilde{X} \in \mathbb{R}^{V \times T}$ denote the surface BOLD time series of a subject measured at cortical locations $\{x_v\}_{v=1}^{V}$. Each row of $\tilde{X}$ is z-scored over the T time points. Functional connectivity among the V locations is given by the spatial Gram matrix $C = \tilde{X}\tilde{X}^\top / T \in \mathbb{R}^{V \times V}$.

We aim to learn z such that the inner product (Eq. 1) between the sampled embeddings approximates the observed functional connectivity: $z(x_i)^\top z(x_j) \approx C_{i,j}$. [1] Directly optimizing this approximation can be computationally heavy as C is often high-dimensional. However, in rs-fMRI studies,

[1] Such an embedding is not unique, since any orthogonal transformation of z preserves these inner products.

$\mathrm{rank}(C) \leq T - 1 \ll V$, which implies C is strongly rank deficient by construction and need not be explicitly formed. Therefore, we can learn z efficiently through a low-rank spectral factorization of $\tilde{X}$. Specifically, we compute the much smaller temporal Gram matrix $G = \tilde{X}^\top \tilde{X}/T \in \mathbb{R}^{T\times T}$, where $G = W\Lambda W^\top$. In other words, C contains inner products between cortical time series, whereas G contains inner products between whole-cortex spatial patterns at pairs of time points. The two matrices share their non-zero spectrum. Writing the singular value decomposition as $\tilde{X}/\sqrt{T} = U\Sigma W^\top$, the eigenvectors W therefore define temporal modes onto which each cortical time series can be projected. With this construction, the d leading modes,

$$\Phi = \tilde{X}W_d/\sqrt{T} = U_d\Sigma_d \in \mathbb{R}^{V\times d}, \tag{3}$$

gives the optimal functional embedding of the V cortical locations. In other words, $\phi_i^\top \phi_j = C_{ij}$, when $d \geq \mathrm{rank}(C)$, while for smaller d, $\Phi\Phi^\top$ is the optimal rank-d approximation of C in Frobenius norm by the Eckart–Young–Mirsky theorem. The d dimensions in Φ are ordered by the variance captured by their corresponding modes, thereby providing a principled truncation of the connectivity representation. The norm $\|\phi_v\|_2$ additionally reflects how much of the original variance at location v is retained by this truncation.

Therefore, the resulting Φ provides the optimal targets for learning the continuous field. We can simply fit z at the observed cortical vertices to reproduce the normalized spectral embeddings:

$$\mathcal{L} = \frac{1}{V}\sum_{v=1}^{V} \ell\left(z(x_v), \hat{\phi}_v\right), \tag{4}$$

where $\hat{}$ denotes ℓ_2 normalization, i.e., $\hat{u} = u/\|u\|_2$, and ℓ denotes the loss between unit vectors.

3.2 Population and Individual Functional Fields

A field fitted independently to each subject cannot be easily compared among individuals for two reasons. First, the temporal modes used to construct Φ for each subject are identifiable only up to an orthogonal transformation, so they are not necessarily expressed in the same basis across subjects. Second, a field estimated from a single rs-fMRI acquisition may capture session-specific variation in addition to reproducible individual organization. To enable cross-subject comparison, we propose to first estimate a population representation that captures shared functional structure and subsequently represent each individual as a constrained adaptation to this population field.

Population functional field. Given functional embeddings $\{\Phi_s\}_{s=1}^S$ from S training subjects, we concatenate them along the feature dimension and compute a joint low-rank factorization,

$$[\Phi_1\ \Phi_2\ \cdots\ \Phi_S] = U\Sigma V^\top, \qquad \Psi = U_{:,1:d}\Sigma_{1:d}/\sqrt{S} \in \mathbb{R}^{V\times d}. \tag{5}$$

This joint factorization provides a shared low-dimensional representation Ψ of the population functional organization. The population field is obtained by fitting an INR to this shared representation,

$$\mathcal{L}_{\mathrm{pop}} = \frac{1}{V}\sum_{v} \left\| z_{\mathrm{pop}}(x_v) - \hat{\Psi}_v \right\|^2. \tag{6}$$

The resulting $z_{\mathrm{pop}}(x)$ provides a shared continuous reference from which subject-specific functional organization can be estimated.

Individualized functional fields. We now compute subject-specific fields by parameterizing each individual as a continuous residual adaptation of the population field, for $x \in \mathbb{S}^2$:

$$z_s(x) = z_{\mathrm{pop}}(x) + \delta_s(x), \qquad \delta_s(x) = \mathrm{MLP}_s(\mathrm{enc}(x)), \tag{7}$$

where δ_s is a subject-specific INR network. Because both components are functions of cortical position, the individualized representation remains continuous over the cortical surface. Unlike Φ_s, z_s in Eq. 7 is anchored on the same population field and therefore comparable across individuals.

The individual field is optimized to reproduce the subject's FC while remaining close to the population representation. Let $\hat{Z}_s$ and Δ_s denote the collection of embeddings $\hat{z}_s(x_v)$ and $\delta_s(x_v)$ over all cortical vertices, we define

$$\mathcal{L}_s = 1 - \mathrm{corr}\left(\hat{Z}_s\hat{Z}_s^\top, C_s\right) + \frac{\lambda}{V}\|\Delta_s\|_F^2, \tag{8}$$

where $\hat{\cdot}$ denotes row-wise ℓ_2 normalization and the correlation is computed over off-diagonal entries. The first term encourages the field to reproduce the subject-specific connectivity C_s, while the second constrains the magnitude of the individual deviation. The parameter λ therefore controls the trade-off between individual specificity and similarity to the shared population representation. To avoid constructing the full vertex-level connectivity matrix, we optimize Eq. 8 stochastically over uniformly sampled cortical locations; implementation details are provided in Appendix A.3.

3.3 Differentiable Spatial Readouts

A continuous functional field changes the role of spatial regions in FC analysis. Given a phenotype of interest, conventional analyses typically test which regions from a predefined parcellation show statistical associations with the phenotype. However, the spatial scale and boundaries of these *a priori* defined regions may not correspond to the functional territories most relevant to that phenotype. With a continuous functional field, we can instead treat the location and geometry of cortical regions as optimizable quantities. This allows us to search continuously over cortical space for regions whose functional connectivity is maximally associated with a downstream phenotype, without restricting the search to the boundaries of a predefined atlas.

Continuous regions. We represent a cortical region by a differentiable membership function $m(x;\vartheta) \in [0,1]$, where $x \in \mathbb{S}^2$ and ϑ parameterizes the region's location, scale, and shape. Given the individualized functional field $z_s(x)$, we obtain a spatial readout of a region by pooling the normalized functional embeddings according to its spatial membership,

$$e_s(\vartheta) = \frac{\int_{\mathbb{S}^2} m(x;\vartheta)\, \hat{z}_s(x)\, d\Omega}{\int_{\mathbb{S}^2} m(x;\vartheta)\, d\Omega}. \tag{9}$$

Critically, because both $z_s(x)$ and $m(x;\vartheta)$ vary continuously over the cortex, the resulting readout $e_s(\vartheta)$ is differentiable with respect to the geometry parameters ϑ.

Region-level functional connectivity. Connectivity between two spatial readouts is defined by the similarity of their pooled functional representations,

$$\rho_s(\vartheta_1, \vartheta_2) = \frac{e_s(\vartheta_1)^\top e_s(\vartheta_2)}{\|e_s(\vartheta_1)\|_2 \|e_s(\vartheta_2)\|_2}. \tag{10}$$

Thus, the same individualized field can be queried with arbitrary region definitions without recomputing the underlying functional representation. Changing the region parameters changes which parts of the continuous field are pooled and, consequently, the resulting region-level FC.

Spatial optimization for brain-phenotype association. We use this differentiability to adapt spatial readouts to a downstream endpoint. Specifically, we jointly optimize the locations and boundary shapes of the two regions to maximize the FC between their FC and a phenotypic variable of interest, y_s. This can be written generally as

$$\max_{\vartheta_1, \vartheta_2} \mathcal{J}\left(\{\rho_s(\vartheta_1, \vartheta_2), y_s\}_{s=1}^N\right), \tag{11}$$

where $\mathcal{J}$ measures the association between region-level connectivity and the phenotype, e.g., the correlation between FC and age.

Region parameterization. In practice, we parameterize each spatial readout as a soft star-shaped region on the cortical sphere, initialized as a geodesic disc and described by its center, radius, and boundary contour. The contour is represented by a low-dimensional Fourier expansion controlling deviations from a circular boundary, and these geometric components can be fixed or optimized for the phenotype of interest. Throughout this optimization, the functional fields remain fixed, such that only their spatial readout is adapted. The formulation extends to P regions by combining their $P(P-1)/2$ pairwise connections through a regularized linear readout. Exact parameterization, optimization procedures, and the multi-region extension are provided in Appendix A.1.

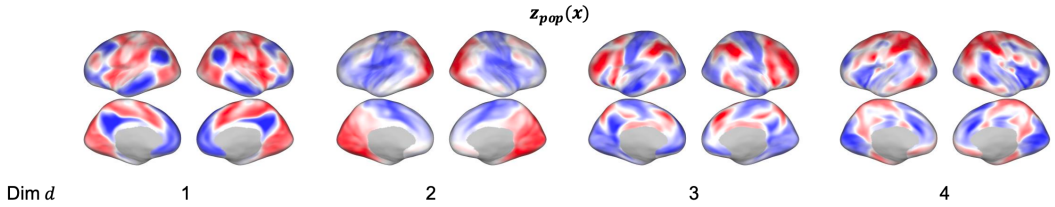


Figure 2: First four dimensions of the learned population functional field $z_{\text{pop}}(x)$.

4 Experiments & Results

In this study, we used resting-state functional MRI data from three independent cohorts: the Human Connectome Project–Development (HCP-D; $n = 625$, age 8–22 years, 337 females) (Somerville et al., 2018), the Philadelphia Neurodevelopmental Cohort (PNC; $n = 1,418$, age 8–23 years, 761 females) (Satterthwaite et al., 2014), and the Healthy Brain Network (HBN; $n = 2,207$, age 5–22 years, 776 females) (Alexander et al., 2017). All HCP-D participants have two rs-fMRI sessions, enabling evaluation across acquisitions. Group-difference and phenotypic analyses were restricted to participants with complete information. Structural and functional MRI data of each individual were preprocessed using the DeepPrep pipeline (Ren et al., 2025), which produces BOLD time series for $V = 81924$ cortical vertices. Additional dataset characteristics, preprocessing steps, and cohort splits are provided in Appendix A.2.

4.1 Learning Population and Individualized Functional Fields

We first examine the dimensionality of functional embeddings d required to represent the observed FC. For each individual, we measure the correlation between C_s and the rank-d approximation $\Phi_s\Phi_s^\top$ derived by the INR models, as well as the variance in C_s retained by the low-rank approximation. This experiment is conducted on HCP-D because it is the only dataset with two rs-fMRI sessions for each individual, allowing us to assess test–retest reliability computed as the correlation of $\Phi_s\Phi_s^\top$ between two runs. Table 1 indicates reconstruction progressively increases with rank and reaches the maximum at $d = 128$, where 99.0% of the Gram spectrum is retained. Test–retest agreement similarly plateaus at $d = 128$. We use $d = 128$ throughout the remaining experiments.

Next, we learn population-level functional fields from HCP-D to characterize shared functional organization. Figure 2 shows the first four dimensions of the population embedding, Appendix Figure 4 compared them alongside four FC gradients (Margulies et al., 2016). The maps exhibit qualitative spatial correspondence, suggesting that the population fields capture aspects of cortical functional organization reflected in FC gradients. Guided by the population fields, we derive individualized fields z_s. To examine the effect of the residual penalty λ on individualization, we run another test-retest analysis, where we fit z_s on the first rs-fMRI session and assess the agreement $\text{corr}(\hat{Z}_s\hat{Z}_s^\top, C_s)$ for both sessions (Table 2). Without residual regularization ($\lambda = 0$), the individualized field nearly perfectly reconstructs C_s used for fitting ($r = 0.991$), but its agreement with the held-out run is lower ($r = 0.525$). As λ increases, z_s is progressively constrained towards the population field, reducing in-sample agreement but improving cross-run reproducibility. Held-out agreement peaks at $\lambda = 1$ and declines again with stronger population constraints. This pattern indicates that unconstrained individualization captures substantial acquisition-specific variation, whereas moderate population anchoring preferentially retains subject-specific organization that generalizes across runs.

Table 1: Effect of embedding rank d in 104 HCP-D subjects with two runs each. Reconstruction is the correlation between rank-d and measured FC, test–retest is the correlation from the two runs, and variance retained is the fraction of the FC spectrum captured by the rank-d representation.

d	4	8	16	32	64	128	256
Reconstruction	0.8782	0.9473	0.9805	0.9938	0.9987	**1.0000**	1.0000
Test–retest	0.4654	0.5021	0.5166	0.5203	0.5220	**0.5223**	0.5223
Variance retained	0.364	0.515	0.665	0.799	0.914	**0.990**	1.000

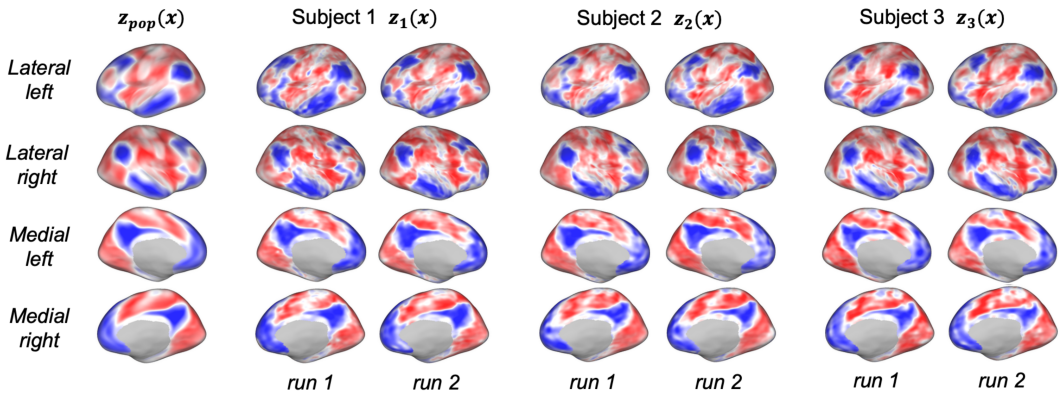


Figure 3: First dimension of the population field $z_{\text{pop}}(x)$ and individualized fields $z_s(x)$ for three representative subjects, estimated independently from two rs-fMRI runs.

We therefore use $\lambda = 1$ in the subsequent experiments. Figure 3 compares the first dimension of the population field with individualized fields from three subjects across two sessions. Individualized maps preserve the population organization while showing finer-scale subject differences that appear consistent across sessions. The Fourier frequencies used by the residual INR network (Eq. 7) have little effect on held-out agreement (Appendix A.3, Table 8).

4.2 Continuous Spatial Readouts Support Brain–Phenotype Mapping

We next examine the utility of individualized functional fields for phenotype association analyses. For each phenotype, a pair of candidate cortical locations is selected on the training subjects, and these regions are optimized to maximize the association between their connectivity and the phenotype. The fitted geometry is then frozen and applied to held-out subjects, and we report the squared t^2 statistic for the brain-phenotype association. We compare against the Schaefer parcellation (Schaefer et al., 2018), a widely used functional atlas, at 100 and 1000 parcels, which span the range of spatial scales commonly used. For each Schaefer parcellation, we select a parcel pair whose FC is most strongly associated with the phenotype on the training subjects and evaluate the corresponding association strength on the held-out subjects. Region selection, optimization, and evaluation of association strength are repeated over random subject splits, and we report medians with interquartile ranges. Full details are given in Appendix A.3.

Age and sex across datasets. We first study whether continuous region optimization can increase the sensitivity of identifying association with age and sex, two phenotypes available in all three cohorts. Across 12 random train–test splits, the optimized regions achieve the highest sex-difference and age correlations across all cohorts (Table 3). Notably, the atlas-based results show no consistent preference for a spatial scale: age favors the coarse parcellation in HCP-D but the fine one in HBN, and the two are close in PNC. While for sex, the fine parcellation leads in PNC, and the two overlap substantially in HCP-D and HBN. These results suggest that using a single arbitrarily defined parcellation is suboptimal for phenotype association analysis. Continuous optimization instead allows

Table 2: Effect of residual penalty λ in 104 HCP-D subjects. In-sample and held-out fit are correlations with FC from the fitted and independent run, respectively, averaged over both run orderings; $\|\delta_s\|$ is the mean norm of the residual $\delta_s(x)$ over cortical locations, measuring how far the individual field departs from the population field.

λ	0	0.1	0.3	0.5	1	3	5
Held-out fit	0.525	0.543	0.561	0.571	**0.583**	0.576	0.559
In-sample fit	0.991	0.979	0.955	0.930	0.876	0.731	0.654
$\|\delta_s\|$	0.867	0.577	0.474	0.414	0.323	0.171	0.111

Table 3: Held-out t^2 for age and sex across three cohorts, reported as median$_{[\mathrm{IQR}]}$ over 12 splits. Schaefer columns use the best training-selected parcel pair at each resolution; *Ours* reports the best optimized region configuration over $r_0 \in \{4, 8, 12\}^\circ$ and $K \in \{2, 4, 6\}$ (Table 9).

Phenotype	Cohort	Schaefer-100	Schaefer-1000	Ours
sex	HCP-D	$9.9_{[5.7,18.0]}$	$9.4_{[5.8,13.1]}$	$\mathbf{15.0}_{[10.5,16.3]}$
	PNC	$8.0_{[5.4,14.3]}$	$14.0_{[10.0,16.3]}$	$\mathbf{18.6}_{[12.5,22.7]}$
	HBN	$4.7_{[2.8,5.9]}$	$3.9_{[1.4,8.1]}$	$\mathbf{17.5}_{[12.6,19.8]}$
age	HCP-D	$67.2_{[60.7,72.5]}$	$31.3_{[20.8,41.1]}$	$\mathbf{78.4}_{[68.5,97.4]}$
	PNC	$40.4_{[37.4,47.4]}$	$36.7_{[28.6,50.8]}$	$\mathbf{84.2}_{[76.8,91.4]}$
	HBN	$51.0_{[44.6,64.8]}$	$67.5_{[61.2,71.8]}$	$\mathbf{116.1}_{[93.2,138.7]}$

Table 4: Held-out t^2 for behavioral and cognitive phenotypes in HCP-D, reported as median$_{[\mathrm{IQR}]}$ over 30 splits. Age and sex are regressed from FC and phenotype. Schaefer columns use the best training-selected parcel pair at each resolution; *Ours* uses optimized regions.

Target	n	Schaefer-100	Schaefer-1000	Ours
CBCL internalizing	373	$0.41_{[0.08,1.03]}$	$0.49_{[0.06,0.95]}$	$0.59_{[0.06,0.97]}$
CBCL externalizing	373	$0.24_{[0.06,0.72]}$	$0.77_{[0.12,2.17]}$	$0.91_{[0.19,1.70]}$
NIHTB Flanker	366	$0.74_{[0.17,1.21]}$	$0.81_{[0.32,2.09]}$	$1.08_{[0.32,2.12]}$
NIHTB fluid comp.	366	$0.44_{[0.11,0.98]}$	$0.32_{[0.05,0.93]}$	$0.73_{[0.21,1.52]}$
NIHTB crystallized	366	$0.80_{[0.46,1.69]}$	$1.16_{[0.23,2.57]}$	$1.35_{[0.41,2.73]}$

the spatial support of FC to adapt to the phenotype and population under study, without requiring its scale or boundaries to be specified *a priori*.

Behavioral associations within HCP-D. We next associate FC with behavioral and cognitive phenotypes in HCP-D (Table 4): CBCL internalizing and externalizing scores ($n = 373$) (Achenbach & Rescorla, 2001), and NIH Toolbox (NIHTB) Flanker, fluid cognition, and crystallized cognition scores ($n = 366$) (Weintraub et al., 2013). Because age and sex are associated with both connectivity and phenotypes, they are treated as covariates and regressed out of each separately before computing the association. Results are summarized as medians across 30 repeated splits, with 35% of subjects held out for testing. The association between the phenotype and the FC between the optimized regions yields a higher median held-out t^2 than the best region pairs in Schaefer-100 and -1000 parcellations for each of the five reported outcomes. The improvement in t^2 ranges from 16% to 66%, including a 33% increase for Flanker and a 66% increase for fluid cognition. Thus, the advantage of continuous spatial optimization extends beyond age and sex to behavioral and cognitive phenotypes with substantially weaker FC associations (Marek et al., 2022).

Extension from a single region pair to multiple regions. The spatial readout also supports joint analysis of multiple connections. For P regions, all $M = P(P-1)/2$ connections are combined using a regularized linear readout (Appendix Eq. 21, Appendix A.3). Both shape and readout weights are fitted on training subjects and frozen for evaluation. In the HCP-D sex experiment (Table 5), increasing from three regions (3 connections) to eight regions (28 connections) raises median held-out t^2 from 13.66 to 25.54 for the optimized contours. The association strength also exceeds both the Schaefer-100 and -1000 comparators at every matched region count. These results demonstrate that continuously optimized regions can serve as building blocks for multivariate FC analysis, rather than being restricted to identifying a single phenotype-associated connection.

Parameterization of region geometry. A region can be parameterized with different degrees of geometric flexibility: optimization may adjust only its location, additionally adjust its size (a disc with variable radius), or further allow its boundary to deform. We compare these parameterizations using the HCP-D age target, optimizing every combination from different initial radius[2] r_0 over 30 repeated splits (35% testing subjects; Table 6). When the radius is not optimized, r_0 remains fixed. For contour models, $K \in \{2, 4, 6\}$ is selected separately at each r_0. Every fitted family improves

[2]Defined as degree in cortical sphere

Table 5: Multi-region HCP-D sex association using the regularized Hotelling readout (Eq. 21). P regions define $P(P-1)/2$ connections. Schaefer uses P parcels from each atlas; *Ours* optimizes region geometry. Values are median$_{[\mathrm{IQR}]}$ held-out t^2 over 30 splits.

P	Edges	Schaefer-100	Schaefer-1000	Ours
3	3	$9.71_{[7.90,14.16]}$	$10.79_{[6.00,16.61]}$	$\mathbf{13.66}_{[9.68,17.64]}$
4	6	$11.19_{[8.43,16.35]}$	$15.26_{[10.38,21.09]}$	$\mathbf{17.78}_{[11.51,22.11]}$
6	15	$16.03_{[12.26,20.79]}$	$20.39_{[14.75,24.72]}$	$\mathbf{22.27}_{[17.89,28.18]}$
8	28	$19.26_{[16.82,23.05]}$	$22.28_{[17.75,29.00]}$	$\mathbf{25.54}_{[18.87,33.83]}$

Table 6: Geometry ablation for HCP-D age. Rows indicate optimized region parameters and columns the initial radius r_0. Values are median$_{[\mathrm{IQR}]}$ held-out t^2 over 30 splits; Range is the difference between the largest and smallest median across r_0. For contour models, $K \in \{2, 4, 6\}$ is selected separately for each r_0.

Components	r_0=4°	8°	12°	Variability in t^2
None (unfitted disc)	$38.88_{[30.6,46.2]}$	$44.83_{[35.4,58.2]}$	$47.36_{[34.1,64.5]}$	8.48
Center	$56.28_{[37.5,68.7]}$	$68.68_{[45.0,81.1]}$	$78.59_{[48.0,87.9]}$	22.31
Radius	$45.67_{[34.9,60.7]}$	$45.68_{[34.9,60.7]}$	$49.27_{[34.9,63.9]}$	3.60
Contour	$55.53_{[35.5,65.6]}$	$71.42_{[46.9,86.1]}$	$76.70_{[51.5,93.4]}$	21.17
Centre + radius	$69.37_{[47.2,82.6]}$	$69.34_{[46.7,82.5]}$	$75.28_{[48.5,89.0]}$	5.94
Centre + contour	$67.02_{[47.4,77.1]}$	$77.57_{[51.7,89.1]}$	$75.36_{[50.0,90.2]}$	10.55
Radius + contour	$73.30_{[49.7,88.8]}$	$72.34_{[49.7,88.6]}$	$73.86_{[48.7,86.3]}$	1.52
Centre + radius + contour	$74.02_{[51.7,87.6]}$	$75.28_{[50.1,89.0]}$	$74.76_{[49.9,89.3]}$	**1.26**

over the unfitted disc, but center- or contour-only optimization is highly sensitive to initialization (e.g., variability in held-out $t^2 > 20$). Optimizing radius reduces this dependence, with joint center, radius, and contour optimization giving the smallest variability across initial radii and the highest worst-case performance. Greater flexibility does not simply improve generalization: training t^2 increases with model capacity, whereas held-out performance does not (Appendix 10).

5 Discussion & Conclusion

We introduced a continuous formulation of functional connectivity that replaces parcellation-defined FC matrices with low-dimensional neural fields over the cortical surface. By combining a shared population field with constrained subject-specific adaptations, the framework captures individualized functional organization while maintaining a common continuous representation across subjects. Beyond representation, their differentiability enables cortical regions to be treated as spatial readouts whose location, scale, and geometry can be optimized for a downstream question. Across three independent cohorts, these optimized readouts consistently yielded stronger held-out associations with age and sex than fixed parcellations, with similar improvements for behavioral and cognitive phenotypes and when extending the formulation to multiple regions.

Our study also has several limitations. First, the current formulation is restricted to the cortical surface and therefore does not model subcortical or cerebellar structures and interactions. Second, individualized fields are estimated from finite resting-state acquisitions and remain limited by the reliability of the underlying FC measurements; indeed, our test–retest experiments show that unconstrained individualization can capture acquisition-specific variation, motivating the population regularization used here. Finally, although we evaluate spatial readouts across multiple cohorts and phenotypes, the present experiments focus on developmental datasets and relatively simple association objectives. Extending evaluation to broader populations, clinical endpoints, and predictive settings will be important to establish the generality of the framework.

More broadly, our results suggest that functional connectivity need not be tied to a particular parcellation or spatial scale. A continuous functional field instead separates the representation of brain function from the spatial regions used to interrogate it, allowing different questions to induce different readouts of the same underlying representation.

AI USE STATEMENT

In this work, we used generative AI tools for improving readability and implementing methods. We have not used generative AI tools for these tasks: generate synthetic data sets, help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, design or provide feedback on research methodology or experiments, assist with translation, clean and reformat dataset, support qualitative and thematic data analysis, interpret results. And the following tasks: assist in the writing of proofs, propose or refine hypotheses, are not applicable to this work. We have reviewed all AI-assisted work. For example, LLM-generated code was verified and tested for correctness. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This study uses data from the publicly available HCP-Development (HCP-D, `https://www.humanconnectome.org/study/hcp-lifespan-development`), Philadelphia Neurodevelopmental Cohort (PNC, `https://www.med.upenn.edu/bbl/philadelphianeurodevelopmentalcohort.html`), and Healthy Brain Network (HBN, `https://fcon_1000.projects.nitrc.org/indi/cmi_healthy_brain_network/`) datasets. HCP-D was approved by the Washington University Institutional Review Board, PNC by the Institutional Review Boards of the University of Pennsylvania and the Children's Hospital of Philadelphia, and HBN by the Chesapeake Institutional Review Board. Informed consent, parental permission, and assent were obtained as appropriate in the respective studies. All data were used in accordance with the corresponding data-use agreements.

REPRODUCIBILITY STATEMENT

The main text provides the methodological and experimental details necessary to understand the proposed framework and reported results. For full reproducibility, the appendix provides detailed descriptions of the datasets and preprocessing (A.2), model architectures, training procedures, hyperparameter selection, spatial readout parameterization, and optimization and evaluation protocols(A.3). Source code and implementation details will be made publicly available upon acceptance.

REFERENCES

Thomas M. Achenbach and Leslie A. Rescorla. *Manual for the ASEBA School-Age Forms & Profiles*. University of Vermont, Research Center for Children, Youth, & Families, Burlington, VT, 2001.

Lindsay M. Alexander, Jasmine Escalera, Lei Ai, Charissa Andreotti, Karina Febre, Alexander Mangone, Natan Vega-Potler, Nicolas Langer, Alexis Alexander, Meagan Kovacs, et al. An open resource for transdiagnostic research in pediatric mental health and learning disorders. *Scientific Data*, 4:170181, 2017. doi: 10.1038/sdata.2017.181.

Bharat Biswal, F. Zerrin Yetkin, Victor M. Haughton, and James S. Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, 34(4):537–541, 1995. doi: 10.1002/mrm.1910340409.

Haomiao Chen, Keith W. Jamison, Mert R. Sabuncu, and Amy Kuceyeski. Beyond grid-locked voxels: Neural response functions for continuous brain encoding. In *International Conference on Learning Representations*, 2026.

Emily S. Finn, Xilin Shen, Dustin Scheinost, Monica D. Rosenberg, Jessica Huang, Marvin M. Chun, Xenophon Papademetris, and R. Todd Constable. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, 18(11):1664–1671, 2015. doi: 10.1038/nn.4135.

Michael D. Fox, Abraham Z. Snyder, Justin L. Vincent, Maurizio Corbetta, David C. Van Essen, and Marcus E. Raichle. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, 102(27):9673–9678, 2005. doi: 10.1073/pnas.0504136102.

Evan M. Gordon, Timothy O. Laumann, Adrian W. Gilmore, Dillan J. Newbold, Deanna J. Greene, Jeffrey J. Berg, Mario Ortega, Catherine Hoyt-Drazen, Caterina Gratton, Haoxin Sun, Jacqueline M. Hampton, Timothy S. Coalson, Annie L. Nguyen, Kathleen B. McDermott, Joshua S. Shimony, Abraham Z. Snyder, Bradley L. Schlaggar, Steven E. Petersen, Steven M. Nelson, and Nico U. F. Dosenbach. Precision functional mapping of individual human brains. *Neuron*, 95(4): 791–807.e7, 2017. doi: 10.1016/j.neuron.2017.07.011.

Andrew Hannum and Mario A. Lopez. Task-optimized brain parcellations reveal latent functional organization for enhanced connectivity-based neuroimaging classification. *NeuroImage*, 325: 121689, 2026. doi: 10.1016/j.neuroimage.2026.121689.

Scott Marek, Brenden Tervo-Clemmens, Finnegan J. Calabro, et al. Reproducible brain-wide association studies require thousands of individuals. *Nature*, 603:654–660, 2022. doi: 10.1038/s41586-022-04492-9.

Daniel S. Margulies, Satrajit S. Ghosh, Alexandros Goulas, Marcel Falkiewicz, Julia M. Huntenburg, Georg Langs, Gleb Bezgin, Simon B. Eickhoff, F. Xavier Castellanos, Michael Petrides, Elizabeth Jefferies, and Jonathan Smallwood. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences*, 113(44):12574–12579, 2016. doi: 10.1073/pnas.1608282113.

Vincent Michel, Alexandre Gramfort, Gaël Varoquaux, Evelyn Eger, Christine Keribin, and Bertrand Thirion. A supervised clustering approach for fMRI-based inference of brain states. *Pattern Recognition*, 45(6):2041–2049, 2012. doi: 10.1016/j.patcog.2011.04.006.

Sophia Mueller, Danhong Wang, Michael D. Fox, B. T. Thomas Yeo, Jorge Sepulcre, Mert R. Sabuncu, Rebecca Shafee, Jimin Lu, and Hesheng Liu. Individual variability in functional connectivity architecture of the human brain. *Neuron*, 77(3):586–595, 2013. doi: 10.1016/j.neuron.2012.12.028.

Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 165–174, 2019.

Jianxun Ren, Ning An, Cong Lin, Youjia Zhang, Zhenyu Sun, Wei Zhang, Shiyi Li, Ning Guo, Weigang Cui, Qingyu Hu, et al. Deepprep: an accelerated, scalable and robust pipeline for neuroimaging preprocessing empowered by deep learning. *Nature Methods*, 22:473–476, 2025. doi: 10.1038/s41592-025-02599-1.

Mehraveh Salehi, Amin Karbasi, Xilin Shen, Dustin Scheinost, and R. Todd Constable. An exemplar-based approach to individualized parcellation reveals the need for sex specific functional networks. *NeuroImage*, 170:54–67, 2018. doi: 10.1016/j.neuroimage.2017.08.068.

Theodore D. Satterthwaite, Mark A. Elliott, Kosha Ruparel, James Loughead, Karthik Prabhakaran, Monica E. Calkins, Ryan Hopson, Chad Jackson, Jack Keefe, Marisa Riley, et al. Neuroimaging of the philadelphia neurodevelopmental cohort. *NeuroImage*, 86:544–553, 2014. doi: 10.1016/j.neuroimage.2013.07.064.

Alexander Schaefer, Ru Kong, Evan M. Gordon, Timothy O. Laumann, Xi-Nian Zuo, Avram J. Holmes, Simon B. Eickhoff, and B. T. Thomas Yeo. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cerebral Cortex*, 28(9):3095–3114, 2018. doi: 10.1093/cercor/bhx179.

Xilin Shen, Emily S. Finn, Dustin Scheinost, Monica D. Rosenberg, Marvin M. Chun, Xenophon Papademetris, and R. Todd Constable. Using connectome-based predictive modeling to predict individual behavior from brain connectivity. *Nature Protocols*, 12(3):506–518, 2017. doi: 10.1038/nprot.2016.178.

Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems*, volume 33, pp. 7462–7473, 2020.

Stephen M. Smith, Diego Vidaurre, Christian F. Beckmann, Matthew F. Glasser, Mark Jenkinson, Karla L. Miller, Thomas E. Nichols, Emma C. Robinson, Gholamreza Salimi-Khorshidi, Mark W. Woolrich, Deanna M. Barch, Kamil Uğurbil, and David C. Van Essen. Functional connectomics from resting-state fMRI. *Trends in Cognitive Sciences*, 17(12):666–682, 2013. doi: 10.1016/j.tics.2013.09.016.

Leah H. Somerville, Susan Y. Bookheimer, Randy L. Buckner, Gregory C. Burgess, Sarah W. Curtiss, Mirella Dapretto, Jennifer S. Elam, Michael S. Gaffrey, Michael P. Harms, Cynthia Hodge, et al. The lifespan human connectome project in development: A large-scale study of brain connectivity development in 5–21 year olds. *NeuroImage*, 183:456–468, 2018. doi: 10.1016/j.neuroimage.2018.08.050.

Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems*, volume 33, pp. 7537–7547, 2020.

Sandra Weintraub, Patricia J. Bauer, Philip D. Zelazo, Kathleen Wallner-Allen, Sureyya S. Dikmen, Robert K. Heaton, et al. NIH Toolbox Cognition Battery (CB): Introduction and Pediatric Data. *Monographs of the Society for Research in Child Development*, 78(4):1–15, 2013. doi: 10.1111/mono.12031.

Wanwan Xu, Selena Wang, Simiao Gao, Xinyuan Tian, Chichun Tan, Xilin Shen, Wenjing Luo, Todd Constable, Tianxi Li, and Yize Zhao. Supervised brain node and network construction under voxel-level functional imaging. *Imaging Neuroscience*, 3, 2025. doi: 10.1162/IMAG.a.56.

B. T. Thomas Yeo, Fenna M. Krienen, Jorge Sepulcre, Mert R. Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L. Roffman, Jordan W. Smoller, Lilla Zöllei, Jonathan R. Polimeni, Bruce Fischl, Hesheng Liu, and Randy L. Buckner. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106(3):1125–1165, 2011. doi: 10.1152/jn.00338.2011.

A Supplementary Methods

A.1 Differentiable Spatial Readout Details

Region parameterization. We parameterize each spatial readout as a soft star-shaped region on the cortical sphere. A region is centered at $c \in \mathbb{S}^2$. For a cortical location x, let θ_x denote its geodesic distance from c and φ_x its azimuth in the tangent plane at c. The boundary radius is parameterized as

$$r(\varphi) = R\left[1 + \alpha \tanh\left(\sum_{k=1}^{K} a_k \cos(k\varphi) + b_k \sin(k\varphi)\right)\right], \tag{12}$$

where R defines the reference radius, $\{a_k, b_k\}_{k=1}^{K}$ control deviations from a circular boundary, and α bounds the magnitude of these deformations, so that the boundary stays within $[(1-\alpha)R, (1+\alpha)R]$ and never crosses zero. Each geometric component is represented by an unconstrained parameter so that it can be optimized with a standard gradient method. The center is $c = u/\|u\|_2$ for $u \in \mathbb{R}^3$, normalized on every forward pass so that it always lies on the sphere. The reference radius is

$$R = r_{\min} + (r_{\max} - r_{\min})\,\sigma(\eta), \tag{13}$$

which keeps R strictly inside $[r_{\min}, r_{\max}]$ for any value of the scalar η. The region parameters are therefore

$$\vartheta = \left(u,\ \eta,\ \{a_k, b_k\}_{k=1}^{K}\right). \tag{14}$$

Any subset of these components can be held fixed at its initial value. Setting all harmonic coefficients to zero recovers a geodesic disc of radius R.

Soft membership. To maintain differentiability with respect to the region parameters, hard region boundaries are replaced by a soft membership function,

$$m(x; \vartheta) = \sigma\left(\frac{r(\varphi_x) - \theta_x}{\tau}\right) h(x), \tag{15}$$

where σ is the logistic sigmoid, τ controls boundary softness, and $h(x)$ is the cortical mask. The temperature τ can be decreased during optimization to progressively sharpen the region boundary. Integrals in Eqs. 9–10 are evaluated numerically over the cortical surface.

Group contrasts. For a binary endpoint, let $\rho_s(\vartheta_1, \vartheta_2)$ denote the connectivity between two candidate regions for subject s. Spatial parameters are optimized using the squared two-sample t statistic,

$$\mathcal{J}_{\text{group}} = \frac{(\bar{\rho}_1 - \bar{\rho}_0)^2}{s^2_{\text{pooled}}\left(\frac{1}{n_0} + \frac{1}{n_1}\right)}, \tag{16}$$

where $\bar{\rho}_0$ and $\bar{\rho}_1$ are the mean connectivity values in the two groups, n_0 and n_1 their sample sizes, and s^2_{pooled} the pooled variance. Squaring the statistic yields an objective independent of the direction of the group difference.

Continuous phenotypes. For a continuous endpoint y_s, we optimize the association between the endpoint and region-level connectivity. This can be expressed through the squared correlation

$$\mathcal{J}_{\text{cont}} = \text{corr}\left(\{\rho_s(\vartheta_1, \vartheta_2)\}_{s=1}^{N}, \{y_s\}_{s=1}^{N}\right)^2. \tag{17}$$

When covariates are included, the corresponding partial association is used after accounting for those covariates.

Multiple spatial readouts. The formulation can be extended from a single connection to a collection of P regions. Let

$$x_s(\Theta) \in \mathbb{R}^M, \qquad M = \frac{P(P-1)}{2}, \tag{18}$$

collect the pairwise region-level connectivity values induced by region parameters $\Theta = \{\vartheta_p\}_{p=1}^{P}$. For a two-group comparison, let

$$\Delta = \mu_1 - \mu_0 \tag{19}$$

denote the difference between the group-average connectivity vectors. We use a regularized within-group covariance

$$S_{\mathrm{reg}} = S + \lambda_S \frac{\mathrm{tr}(S)}{M} I. \tag{20}$$

The corresponding regularized Hotelling linear readout is

$$w^* = S_{\mathrm{reg}}^{-1} \Delta, \tag{21}$$

yielding the objective

$$\mathcal{J}_{\mathrm{multi}} = \Delta^\top S_{\mathrm{reg}}^{-1} \Delta. \tag{22}$$

This allows the spatial parameters to be optimized according to a multivariate pattern of connectivity rather than a single region pair.

For the Schaefer baselines in the multi-region experiment, the P parcels are selected greedily on the training subjects. All $P(P-1)/2$ parcel-pair connections are ranked by the absolute value of their t statistic for the phenotype, and the ranked list is traversed from the top, adding each connection's two parcels to the selected set if not already present, until the set contains P distinct parcels. All $P(P-1)/2$ connections among the selected parcels are then combined with the same regularized Hotelling readout used for the optimized regions.

Optimization and evaluation. Spatial optimization is performed using only the training subjects. Regions are initialized from predefined cortical locations as geodesic discs, with the deformation coefficients initialized to zero. Region parameters are then optimized through the differentiable functional-field readout while the population and individualized fields remain fixed. After optimization, the learned region parameters are frozen and evaluated on held-out subjects without further spatial selection or fitting. Additional optimization hyperparameters and initialization details are reported in Appendix A.3.

A.2 Dataset and Preprocessing Details

We use three developmental resting-state fMRI cohorts: the Human Connectome Project in Development (HCP-D), the Philadelphia Neurodevelopmental Cohort (PNC), and the Healthy Brain Network (HBN). HCP-D is the discovery cohort for every analysis in the main text; PNC and HBN are used for replication. Table 7 summarizes them. The three differ in the properties that matter most for connectivity estimation: HBN spans a younger age range and four acquisition sites, PNC provides a single short run per subject, and only HCP-D provides two runs for every subject. We conduct our primary experiments on HCP-D resting-state fMRI. Separate rs-fMRI runs are used to evaluate whether individualized functional fields capture reproducible subject-specific organization beyond the acquisition used for fitting.

Table 7: Cohort characteristics. n is the number of subjects for whom individual fields were fitted.

	HCP-D	PNC	HBN
n	625	1,418	2,207
Age, years (mean ± SD)	14.6 ± 3.9	15.3 ± 3.6	10.9 ± 3.4
Age range	8.1–21.9	8.2–23.1	5.0–21.9
Sex (F / M / other)	337/288/0	761/657/0	776/1,358/73
Acquisition sites	1	1	4
Frames per run × runs	478×2	124×1	$375 \times (1 \text{ or } 2)$
Subjects with two runs	625	0	1,674
Median FD (mm)	0.146	0.087	0.175

Preprocessing and Postprocessing Structural and functional MRI data were preprocessed using the DeepPrep pipeline (Ren et al., 2025). Postprocessing included temporal band-pass filtering (0.01–0.08 Hz), regression of nuisance covariates, including head-motion parameters and global signal, spatial smoothing with a 6-mm full-width-at-half-maximum kernel, and projection to the fsaverage6 surface (40,962 vertices per hemisphere; 81,924 vertices in total). After removal of the medial wall, 74,947 cortical vertices were retained for analysis.

A.3 Implementation Details

Coordinate encoding and network architectures. We use separate coordinate networks for the left and right hemispheres, with embedding dimension $d = 128$. For spherical coordinate $x = (x_1, x_2)$ and a set of frequencies $\mathcal{F}$, the coordinate encoding is

$$\mathrm{enc}_{\mathcal{F}}(x) = \big[\sin(mx_1), \cos(mx_1), \sin(mx_2), \cos(mx_2)\big]_{m\in\mathcal{F}}. \tag{23}$$

The population field uses $\mathcal{F} = \{1\}$, giving four input features. Its spatial backbone contains six hidden layers of width 512 with SiLU activations, followed by a 512-to-512 linear projection and a 512-to-128 output head. The output is ℓ_2-normalized at every location. Only the spatial backbone and output head are trained; the population field has no age or phenotype conditioning. Each subject-specific residual network δ_s has four hidden layers of width 256 with SiLU activations and a linear 256-to-128 output head. It uses $\mathcal{F} = \{1, 2, 4, 8, 16, 32\}$, giving 24 input features. The output-head weights and biases are initialized to zero, so $\delta_s(x) = 0$ at initialization and each fit starts from the population field.

Population-field fitting. The population field is fitted once using the HCP-D subjects described in Appendix A.2 . We regress its output directly onto the row-normalized population embedding $\hat{\Psi}$ using Eq. 6. Direct regression preserves the coordinate basis of the population factorization. We optimize the two hemisphere networks jointly with Adam for $9{,}000$ steps, using an initial learning rate of 10^{-3} and cosine annealing to 10^{-4}. Each step samples $4{,}000$ vertices uniformly without replacement from the $74{,}947$ cortical vertices, with a fresh sample at every step. Network initialization and vertex sampling use seed 0. The resulting population field is frozen for all subsequent individual fits in HCP-D, PNC, and HBN.

Individual-field fitting. For a sampled vertex set $\mathcal{B}$, let $\hat{Z}_{s,\mathcal{B}}$ collect the normalized individual embeddings and $\Phi_{s,\mathcal{B}}$ the corresponding spectral embeddings from Eq. 3. The implemented objective is

$$\mathcal{L}_{s,\mathcal{B}} = 1 - \mathrm{corr}_{\mathrm{off}}\left(\hat{Z}_{s,\mathcal{B}}\hat{Z}_{s,\mathcal{B}}^{\top}, \Phi_{s,\mathcal{B}}\Phi_{s,\mathcal{B}}^{\top}\right) + \frac{\lambda}{|\mathcal{B}|}\sum_{v\in\mathcal{B}}\|\delta_s(x_v)\|_2^2. \tag{24}$$

Here, $\mathrm{corr}_{\mathrm{off}}$ is computed by extracting the off-diagonal entries of each sampled Gram matrix, subtracting their mean, normalizing each resulting vector by its ℓ_2 norm, and taking their inner product; 10^{-6} is added to each norm for numerical stability. The spectral target is used without additional row normalization. The residual penalty averages the squared residual norm over sampled vertices, with the sum over embedding dimensions taken inside the norm. We use $\lambda = 1$ and Adam for $4{,}000$ steps, with the learning rate cosine-annealed from 10^{-3} to 10^{-4}. Each step samples $2{,}000$ cortical vertices uniformly without replacement across both hemispheres; there is no separate hemisphere balancing or importance weighting. Both initialization and sampling are reset to seed 0 for every subject and acquisition, giving the same initial weights and sequence of sampled vertex blocks across fits.

Acquisitions used for different cohorts. For HCP-D, we fit separate residual fields to run 1, run 2, and their concatenation. The single-run fits support cross-run evaluation, while the concatenated fit support all other analysis. PNC fields are fitted to the single available run, and HBN fields use all available data from one or two runs. The residual architecture, fitting schedule, and λ are identical across cohorts. The Gram-matching objective is invariant to orthogonal changes of basis in the subject's spectral embedding, allowing all three cohorts to use the same frozen HCP-D population field.

Hyperparameter selection. Embedding rank, residual regularization, and encoding frequencies were examined using 104 HCP-D subjects with two acquisitions each (Tables 1–8). We use $d = 128$, the smallest tested rank at which reconstruction correlation is approximately one and cross-run agreement reaches its observed plateau. This is an empirical sufficiency criterion for the evaluated data. For regularization, we evaluate $\lambda \in \{0, 0.1, 0.3, 0.5, 1, 3, 5\}$ by fitting z_s on each run and measuring agreement between reconstructed and true FC from the other run, averaging over both fit/evaluate directions. The highest mean held-out agreement occurs at $\lambda = 1$, which is then fixed across all cohorts and downstream analyses. This selection uses no phenotype labels and targets. At

fixed $\lambda = 1$, the six tested frequency encodings differ by at most 0.0004 in mean held-out agreement (Table 8), in-sample agreement meanwhile rises monotonically with the number of encoding dimensions, from 0.8573 to 0.8739, so the richer encodings do fit better for the in-sample run without generalizing any better to the other run. We retain $\{1, 2, 4, 8, 16, 32\}$ as the common residual encoding because the choice is immaterial to held-out agreement, not because it outperforms the alternatives.

Table 8: Effect of the Fourier frequencies used by the residual encoding at $\lambda = 1$. Results are computed on the 104 subjects with two rs-fMRI.

m in $\{2^k\}_{k=0}^{m}$	Dim.	In-sample	Held-out
3	16	0.8573	0.5823
4	20	0.8629	0.5825
5	24	0.8687	0.5825
6	28	0.8719	0.5826
7	32	0.8739	0.5827
Dense $\{1, \ldots, 32\}$	40	0.8670	0.5823

Region initialization and optimization. We use the soft geodesic regions defined in Appendix A.1, with boundary-modulation bound $\alpha = 0.8$ in Eq. 12. The sweeps evaluate initial radius $r_0 \in \{4, 8, 12\}^\circ$ and contour orders $K \in \{2, 4, 6\}$. Coefficients are initialized to zero. We use $r_{\min} = 0.5^\circ$ and $r_{\max} = \max(20^\circ, r_0 + 8^\circ)$, which leaves the fitted radius room to grow beyond its initial value. We use Adam with a learning rate of 0.02 and clip the gradient norm to 1. The temperature is annealed through $\tau \in \{1.5, 0.9, 0.5\}^\circ$.

Spatial selection and held-out evaluation. The sweeps use $12/30$ repeated training/held-out subject splits per cohort, with 35% of subjects held out in each split. Within each training split, the candidate set consists of all parcel pairs of the Schaefer-1000 parcellation, with connectivity computed from the individualized fields pooled over each parcel. The pair with the largest absolute training-set t statistic is selected, and the two regions are initialized as geodesic discs centered on the centroids of the selected parcels. This selection is repeated independently for every split. Region geometry is then optimized on the training subjects while all functional fields remain frozen, and the fitted geometry is applied to the held-out subjects without refitting.

Multiple-region readouts. For P regions, we use all $M = P(P-1)/2$ pairwise connections and the regularized covariance from Appendix A.1, with $\lambda_S = 0.1$ fixed throughout. Both the region geometry and the linear readout $w^* = S_{\text{reg}}^{-1}\Delta$ are fitted on training subjects. At evaluation, both are frozen, and group separation is measured by the ordinary two-sample t^2 statistic on the scalar ${w^*}^\top x_s(\Theta)$ in held-out subjects.

Numerical evaluation of spatial readouts. Evaluating field during optimization requires the individual field at locations inside the current region, and because the region moves at every gradient step, querying the field networks directly would require a new forward pass at a new set of points on every step. We avoid this by exploiting the fact that, within a split, each region remains near its initial position. Before optimization, each subject's field is evaluated once on a dense polar chart centered at each initial region center, with 124 samples of geodesic distance $\theta \in [0^\circ, 38^\circ]$ and 180 samples of azimuth $\varphi \in [0^\circ, 360^\circ)$, giving approximately $22{,}000$ locations per region. During optimization, the field at any query point is obtained by bilinear interpolation between the four surrounding chart nodes, periodic in φ. The interpolation is differentiable almost everywhere, so gradients with respect to the region parameters pass through it as they would through the network, while the per-step cost no longer depends on network evaluation. The chart radius is chosen so that each region's support remains within it, and charts are constructed once per split and reused across initial radius and contour orders.

B SUPPLEMENTARY RESULTS

B.1 MULTI-DATASET FULL GRID RESULTS

Table 9 reports the complete radius–contour grid underlying the age and sex results in the main text.

Table 9: Full age and sex results across cohorts. Entries are median held-out t^2 [IQR] over 12 splits for initial radius r_0 and contour order K. Schaefer rows show the corresponding atlas-based baselines; bold marks the largest median within each cohort–target block.

	r_0	$K = 2$	$K = 4$	$K = 6$
Sex				
HCP-D	4°	$13.3_{[10.4,16.1]}$	$13.4_{[9.3,15.8]}$	$13.9_{[9.7,14.6]}$
	8°	$12.8_{[10.3,15.5]}$	$13.1_{[7.8,15.2]}$	$13.6_{[9.8,15.5]}$
	12°	$\mathbf{15.0}_{[10.5,16.3]}$	$14.0_{[7.2,17.5]}$	$13.9_{[7.6,16.0]}$
	Sch-100		$9.9_{[5.7,18.0]}$	
	Sch-1000		$9.4_{[5.8,13.1]}$	
PNC	4°	$17.8_{[14.5,21.6]}$	$17.7_{[14.2,22.0]}$	$18.0_{[14.0,20.9]}$
	8°	$16.0_{[14.7,21.5]}$	$16.2_{[14.0,20.5]}$	$16.6_{[13.9,19.4]}$
	12°	$17.9_{[13.4,21.3]}$	$17.8_{[13.9,22.9]}$	$\mathbf{18.6}_{[12.5,22.7]}$
	Sch-100		$8.0_{[5.4,14.3]}$	
	Sch-1000		$14.0_{[10.0,16.3]}$	
HBN	4°	$15.1_{[10.0,18.2]}$	$15.9_{[12.3,18.2]}$	$15.6_{[11.5,17.6]}$
	8°	$16.5_{[12.5,17.2]}$	$16.3_{[12.6,19.1]}$	$16.1_{[12.9,18.9]}$
	12°	$\mathbf{17.5}_{[12.6,19.8]}$	$17.4_{[12.9,19.8]}$	$16.1_{[13.0,19.5]}$
	Sch-100		$4.7_{[2.8,5.9]}$	
	Sch-1000		$3.9_{[1.4,8.1]}$	
Age				
HCP-D	4°	$74.7_{[62.1,90.8]}$	$75.3_{[68.7,87.9]}$	$76.4_{[67.7,95.1]}$
	8°	$76.1_{[66.8,90.7]}$	$74.4_{[69.0,98.0]}$	$75.5_{[68.7,95.2]}$
	12°	$76.2_{[66.5,89.3]}$	$\mathbf{78.4}_{[68.5,97.4]}$	$78.4_{[66.6,95.1]}$
	Sch-100		$67.2_{[60.7,72.5]}$	
	Sch-1000		$31.3_{[20.8,41.1]}$	
PNC	4°	$76.1_{[67.8,91.7]}$	$77.2_{[72.0,89.8]}$	$77.2_{[71.9,91.4]}$
	8°	$78.5_{[71.5,92.1]}$	$80.4_{[73.7,89.5]}$	$78.6_{[75.2,94.5]}$
	12°	$81.6_{[75.0,91.6]}$	$83.6_{[74.8,92.8]}$	$\mathbf{84.2}_{[76.8,91.4]}$
	Sch-100		$40.4_{[37.4,47.4]}$	
	Sch-1000		$36.7_{[28.6,50.8]}$	
HBN	4°	$105.6_{[93.2,136.6]}$	$111.9_{[90.2,128.8]}$	$109.6_{[94.9,124.0]}$
	8°	$100.1_{[93.4,132.3]}$	$\mathbf{116.1}_{[93.2,138.7]}$	$114.5_{[91.4,131.6]}$
	12°	$110.9_{[96.1,125.3]}$	$110.8_{[97.0,127.0]}$	$110.9_{[90.7,113.1]}$
	Sch-100		$51.0_{[44.6,64.8]}$	
	Sch-1000		$67.5_{[61.2,71.8]}$	

Table 10 reports training t^2 for every cell of Table 6, and Table 11 the harmonic order selected at each radius. Training t^2 generally rises as more geometric parameters are fitted, whereas held-out performance does not follow the same pattern. At 8° and 12°, the family with the highest training value is not the best on held-out subjects: the full model leads in training at 8° (188.11) but is outperformed on held-out data by fitting the centre and contour, and at 12° contour-only leads in training (181.10) while fitting the centre alone gives the highest held-out value (78.59) from a training value of only 159.73.

For a descriptive comparison of geometry families, Table 11 reports the contour order that attains the largest median held-out t^2 at each starting radius.

Table 10: Training t^2 for the geometry ablation in Table 6, reported as median [IQR] over 30 splits. Bold marks the largest median at each starting radius.

Components	r_0=4°	8°	12°
None (unfitted disc)	$108.40_{[101.7,117.7]}$	$121.08_{[101.8,134.0]}$	$115.76_{[93.4,128.0]}$
Centre	$135.19_{[124.4,145.5]}$	$150.32_{[133.2,160.8]}$	$159.73_{[142.7,168.9]}$
Radius	$125.43_{[109.6,138.4]}$	$125.46_{[109.1,138.4]}$	$130.29_{[112.6,138.6]}$
Centre + radius	$159.77_{[141.9,171.1]}$	$159.96_{[142.1,171.2]}$	$166.67_{[146.5,178.3]}$
Contour	$137.62_{[125.8,150.4]}$	$172.56_{[155.3,191.2]}$	$181.10_{[166.8,199.0]}$
Centre + contour	$158.20_{[140.7,171.9]}$	$186.92_{[166.3,200.2]}$	$179.58_{[161.4,193.4]}$
Radius + contour	$182.08_{[165.6,199.9]}$	$182.96_{[166.0,199.9]}$	$167.21_{[152.0,185.8]}$
Centre + radius + contour	$\mathbf{186.35}_{[165.9,200.3]}$	$\mathbf{188.11}_{[168.6,204.0]}$	$178.73_{[161.7,194.4]}$

Table 11: Contour order K selected for each geometry family and starting radius, using the largest held-out median over $K \in \{2, 4, 6\}$.

Components	r_0=4°	8°	12°
Contour	2	6	6
Centre + contour	6	6	2
Radius + contour	6	6	2
Centre + radius + contour	4	6	2

B.2 Supplementary Figures

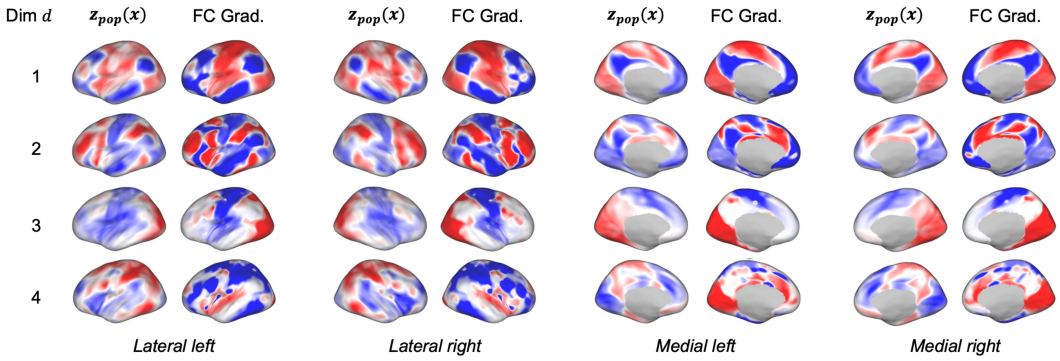


Figure 4: First four dimensions of the learned population functional field $z_{\mathrm{pop}}(x)$ compared with the first four conventional FC gradients across cortical views.

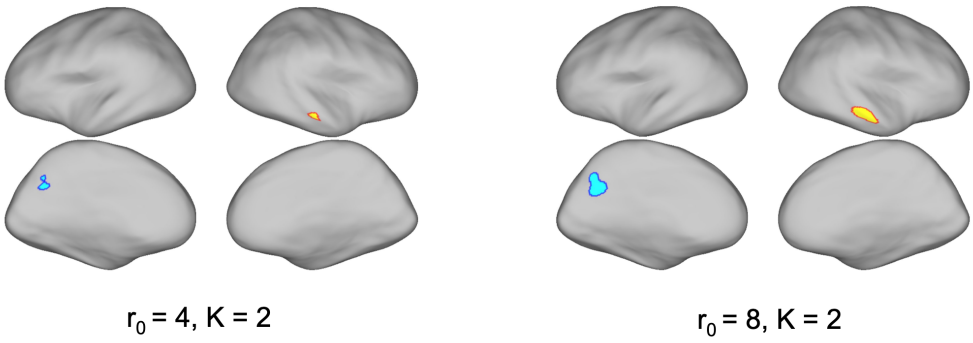


Figure 5: Example of contour-optimized region pairs. The two regions of the pair are shown in yellow and cyan, with outlines marking their boundaries. Both panels use contour order $K = 2$ and differ only in the starting radius, $r_0 = 4^\circ$ (left) and $r_0 = 8^\circ$ (right). Regions are initialized as geodesic discs, and only their boundary contours are optimized.