

Supplementary Information: Charting Higher-Order Models of Brain Function Beyond Pairwise Interactions

Andrea Santoro^{1,*,†}, Matteo Neri^{2,*}, Simone Poetto^{3,4},
Davide Orsenigo⁵, Matteo Diano⁵, Marilyn Gatica^{6,7,‡}, Giovanni Petri^{7,8,‡,§}

¹ ISI Foundation, 10123, Torino, Italy

²Institut de Neurosciences de la Timone UMR 7289, Aix Marseille Université, CNRS, 13005, Marseille, France

³Nicolaus Copernicus University, Toruń, Poland

⁴Intesa Sanpaolo A.I. Research, Torino, Italy

⁵Department of Psychology and Neuroscience Institute of Turin, University of Torino, Torino, Italy

⁶ Centre for Apprenticeships, Northeastern University London, London, E1W 1LP, UK

⁷ NPLab, Network Science Institute, Northeastern University London, London, E1W 1LP, UK

⁸Department of Physics, Northeastern University, Boston, USA

*Equal contribution

‡Joint contribution

† andrea.santoro@isi.it

§giovanni.petri@nulondon.ac.uk

July 10, 2026




Figure S1: **Temporal evolution of research interest in higher-order interactions (HOIs).** The bar charts display the annual volume of scholarly publications indexed by Google Scholar between 1980 and 2025. **(a)** Total number of search results for the specific key phrase “Higher-order interactions”. **(b)** Number of publications identified using the search query “Higher-order Interactions” AND “Neuroscience”. **(c)** Number of publications identified using the search query “Higher-order Interactions” AND “Neuroscience” AND “Brain Connectivity”. The data illustrate a marked increase in relevant literature, particularly over the last decade, with an increasing usage of HOIs as novel brain connectivity over the last 5 years.




Figure S2: Schematic illustration of the comparison of the HOI framework for analyzing fMRI data. (a) fMRI data from the Human Connectome Project (HCP), comprising resting-state and seven task conditions, were parcellated using the Schaefer 100 (cortical) and Tian 16 (subcortical) atlases ($N = 116$). Regional time series were extracted to compute a pairwise baseline (Functional Connectivity, FC) and two categories of HOIs: information theory (IT) metrics (PhiID, PED, O-information, distinguishing between redundancy and synergy) and topological data analysis (TDA) metrics (static scaffolds including Frequency/Persistence, and dynamic metrics like TScaffold and triangles). (b) An edge-level similarity matrix was constructed to classify the computed metrics, revealing a dichotomy in the metric space. When projected at the nodal level, the ‘Redundancy’ cluster (including FC, Redundancy-based IT measures, and triangles) spatially aligns with metabolic gradients and unimodal cortices (visual/somatomotor areas). Conversely, the ‘synergy/topology’ cluster (including synergy-based IT measures and scaffolds) spatially aligns with neurotransmitter receptor density maps and transmodal association cortices. (c) The metrics were evaluated on two distinct prediction tasks: brain fingerprinting (identifying individuals, representing stable “traits”) and task decoding (classifying cognitive states). Results demonstrate a trade-off: HOIs and topological markers maximize individual identifiability (traits), whereas FC and redundancy-dominated measures perform optimally for state decoding. (d) (Left) A ternary plot illustrates the decomposition of FC into redundant, synergistic, and topological percentages, showing that the balance of these components varies a lot for different task conditions. (Right) Partial least squares correlation (PLSC) analysis was applied to link brain features with cognitive scores (e.g., Fluid Intelligence, Memory). The analysis reveals that topological measures explain the most variance in complex behavior compared to other metrics. Brain plots created and modified using the package of Ref. [1].

Nodal projection and inter-subject variability

To complement the results presented in the main text, Supplementary Fig. S1 shows the distribution of average nodal projection ranks across subjects for each of the seven resting-state networks (RSNs) defined by the functional parcellation of Schaefer et al. [2], along with 16 Subcortical areas (SC) of the Tian atlas [3]. These RSNs include: Visual (VIS), Somatomotor (SM), Dorsal Attention (DA), Ventral




Figure S3: **Comparison of nodal projection across different resting state networks.** This panel displays the distribution across subjects of the average nodal strength for each resting-state network (RSN) and subcortical areas — Visual (VIS), Somatomotor (SM), Dorsal Attention (DA), Ventral Attention (VA), Limbic (L), Frontoparietal (FP), Default Mode Network (DMN) and subcortical areas (SC) — for all metrics under investigation. As Pscaff and Fscaff yield highly similar results, only Fscaff is presented for clarity. Notably, nodal projections reveal opposing trends for redundancies and synergies: synergies exhibit stronger nodal strength in multimodal brain regions, whereas redundancies dominate in unimodal, sensorimotor areas. Box plots show the median, interquartile range (IQR), and $1.5 \times \text{IQR}$ whiskers.

Attention (VA), Limbic (L), Frontoparietal (FP), and Default Mode Network (DMN). For each subject, nodal projection values were ranked across all regions from lowest to highest. The ranks of regions belonging to the same RSN were then averaged, yielding a mean rank per RSN per subject. Finally, the across-subject distribution of these average ranks was plotted for each RSN. This analysis was repeated for all the metrics and the results reveal a consistent pattern: unimodal sensorimotor networks tend to rank higher for redundancy-based metrics, whereas multimodal associative networks tend to rank higher for synergy-based metrics. Moreover, we observed that the subject variability is not significantly different between the different metrics.

Network differences between higher-order descriptors

To further understand how functional interactions distribute across the brain at the edge level, we examined the relative strength of connections occurring within versus between RSNs. This comparison sheds light on whether certain metrics preferentially capture cohesive modules or more distributed connectivity patterns. As illustrated in Fig. S4, redundancy-based metrics and scaffold descriptors exhibit significantly stronger within-network connections, consistent with known modular organization of the brain FC [4]. Conversely, synergy-based metrics, a part O-Syn, exhibit stronger connections between regions belonging to different RSNs, capturing patterns of information integration across distinct functional modules of the brain (Fscaff: $t = 17.50$, $p = 2.56 \times 10^{-59}$; Pscaff: $t = 16.17$, $p = 7.10 \times 10^{-52}$; PED Syn: $t = -17.61$, $p = 1.21 \times 10^{-60}$; PhiID Syn: $t = -24.13$, $p = 9.11 \times 10^{-101}$; TScaffold: $t = 31.89$, $p = 1.95 \times 10^{-153}$; Triangles: $t = 10.77$, $p = 1.05 \times 10^{-25}$; O-Red: $t = 17.59$, $p = 7.56 \times 10^{-61}$; PED Red: $t = 17.78$, $p = 1.38 \times 10^{-61}$; PhiID Red: $t = 25.16$, $p = 1.68 \times 10^{-107}$; FC: $t = 25.10$, $p = 3.87 \times 10^{-109}$; O-Syn: $t = 1.87$, $p = 0.062$). These results highlight a fundamental distinction in the network structure related to different metrics. Notably, our results are consistent with previous observations for PhiID redundancy and synergy [5], while also extending these trends to a wider set of higher-order descriptors.

Shared spatial structures among metrics

One of the central aims of this study is to determine whether different metrics give rise to similar spatial patterns across the brain. While the main text focuses on edge-level projections—given their capacity to retain more fine-grained information—here we turn to nodal projections to investigate inter-metric similarity at the regional level. We computed pairwise Spearman correlations between nodal projections




Figure S4: **Comparison of Connection Strengths Within and Between Resting-State Networks.** This panel shows the distribution of link strengths connecting pairs of brain regions within the same resting-state network (RSN) versus those connecting regions across different RSNs. Results are shown for all the metrics under study, since Pscaff and Fscaff produce highly similar results, only Fscaff is shown for visualization purposes. T-tests comparing each pair of distributions were performed, consistently yielding p -values < 0.001 for all metrics except for O-Syn. All metrics showed positive t -statistics — indicating stronger within-network connections — except for the synergistic metrics. Box plots show the median, interquartile range (IQR), and $1.5 \times \text{IQR}$ whiskers.




Figure S5: **Clusters at the nodal level.** Similarity between the nodal projection of all measures, quantified with Spearman correlation, reveals clusters of methods that share spatial organization. **(a-b)** Complete-linkage hierarchical clustering using the distance metric $\sqrt{2(1 - \rho)}$ yields a partition that differs from the one obtained from edge-level projections (Fig. 2c). Here, Tscalfold groups with the synergistic metrics; metric labels retain the same colors as in Fig. 2. **(c)** Structure–function as in Fig. 2d but computed at the nodal level. At this scale, two clusters emerge: (i) synergistic and temporal scaffold nodal projections, which anticorrelate with both structural connectivity strength (SC strength) and functional connectivity strength (FC strength), and (ii) redundancies and static scaffold measures, which correlate with both SC and FC strength.

and applied hierarchical clustering to reveal groups of metrics with similar spatial distributions. As shown in Fig. S5, three well-defined clusters emerge, corresponding to redundancy-based (FC, O-Red, PhiID Red, PED Red, Triangles), topological (Pscalf and Fscalf), and synergy-based descriptors (O-Syn, PhiID Syn, PED Syn, Tscalfold). Notably, Tscalfold, which clusters with Pscalf and Fscalf at the edge level, aligns more closely with synergy-based metrics at the nodal level. This shift underscores how the level of analysis (edge vs. node) can significantly affect metric interpretation and the relationships among descriptors.

Divergence between edge and nodal projections

To directly examine how metrics diverge across projection levels, we compared their similarity properties at the edge- and node-level. Specifically, we computed the difference between the correlation matrices obtained at each level, capturing how inter-metric relationships shift depending on the projection. In parallel, we assessed the stability of each metric’s alignment with FC and SC by comparing correlation values across edge- and node-level projections. Metrics exhibiting minimal change across levels were characterized by similar correlation profiles, while others showed substantial deviation—indicating sensitivity to the level of analysis. Among these, Tscalfold displayed the most pronounced divergence, with marked shifts in both clustering and connectivity correlations. These results reveal that information theoretic metrics and triangles behave quite consistently across levels, whereas the scaffolds, in particular Tscalfold, reflect fundamentally different network properties when evaluated at the edge versus node level projection (Fig. S6). These findings emphasize that edge- and node-level projections offer complementary perspectives on brain network organization. Furthermore, they highlight the unique role of Tscalfold as a bridge between redundancy and synergy: while it shares an edge-level correlation structure with other topological descriptors, its nodal projection aligns more closely with synergy-based brain topographies.





Figure S6: **Comparison of Edge- and Node-level Analyses.** Panel (a) shows the point-wise difference between the correlation matrices of the metrics computed from the edge-level projection and those computed from the node-level projection. This matrix highlights that Tscaffold is the metric whose behavior changes most notably, as it shifts cluster when transitioning from edge- to node-level analysis. In panels (b) and (c), the y -axis represents the correlation with FC (panel b) and SC (panel c) computed at the edge level, while the x -axis shows the correlation with FC (panel b) and SC (panel c) at the nodal level. The greater a metric's distance from the bisectrix, the greater the difference in its behavior between the two projection levels. Panels (b) and (c) are consistent with what observed in panel (a), confirming the scaffold as the metrics with the strongest variability across projection levels.

This dual behavior suggests that Tscaffold may serve as an integrative descriptor, capturing relevant correlation structure at the edge level, while involving similar topography as synergy.

Inter-subjects variability

As described in the main text, to further investigate the relationships between different methods at the nodal level, we employed a range of brain topographies capturing distinct organizational hierarchies, as well as the spatial distribution of various receptors, neurotransmitters, and metabolites. Prior work has shown that gradients derived from conjugate metrics—such as redundancy and synergy—relate to known functional hierarchies. We systematically extended this analysis by incorporating a wider array of redundancy- and synergy-based descriptors. In particular, we focused on the relationship between the nodal projection of the metrics and the sensorimotor–association (S-A) axis proposed by Sydnor et al. [6], which captures a core gradient of cortical functional organization. To complement the results reported in the main text, we computed this correlation at the subject level to test the inter-subject variability of our results. For each metric, we computed the correlation between its nodal projection and the S-A axis on a subject-by-subject basis, resulting in a distribution of correlation values across subjects. This approach enabled us to quantify inter-individual variability in alignment with this functional gradient. A one-sample t -test against zero was used to determine whether the observed correlations were consistently positive or negative at the group level. To control for multiple comparisons, Bonferroni correction was applied across all metrics (Fig. S7a). We applied the same approach to assess how inter-subject variability affects correlations between metric-derived nodal projections and cortical maps of receptor, neurotransmitter, and metabolites density maps. For each map and each metric, subject-level correlations were computed

(a)



(b)




Figure S7: **Subject variability.** (a) Raincloud distribution of Spearman correlations between each subject's nodal projection and the Sensorimotor–Association axis defined by Sydnor et al. [6]. Different colors represent different informational or topological metrics. Asterisks denote statistically significant values after Bonferroni correction. Box plots show the median, interquartile range (IQR), and $1.5 \times \text{IQR}$ whiskers. (b) Same plot as in Fig. 3b, with the addition that gray squares indicate correlations that were not significant at the population level (i.e. $p \leq 0.05$ using a two-sided spin-permutation test with 1000 spatial permutations to account for cortical spatial autocorrelation).

and tested for consistency across the population using a one-sample t -test. Results that did not reach significance after Bonferroni correction are shown in grey in Fig. S7b, highlighting variability in how reliably different metrics capture molecular and metabolic spatial gradients across individuals.




Figure S8: **Node-wise HOI-FC relationship.** Following Vázquez-Rodríguez et al. [7], local structure–function coupling was estimated using a multilinear regression model fit separately for each node i . The dependent variable was the functional connectivity FC_{ij} between node i and all other nodes $j \neq i$, and the independent variables were the 10 HOI metrics for each pair (i, j) . Model parameters were estimated via ordinary least squares. Goodness of fit for each node was quantified by R_i^2 between observed and predicted FC_{ij} .




Figure S9: **Behavioral PLSC scores** . Significant behavioral loadings from the Partial Least Squares Correlation (PLSC) analysis linking brain-level metrics to behavioral dimensions. Each panel shows the loading values of behavioral measures for a given brain feature (e.g., O-Red, PhiID Syn, Scaffold, Functional Connectivity), with colored bars indicating statistically significant loadings. Error bars reflect 95% confidence intervals (CI) from bootstrapping. Non-significant loadings (CI crosses zero) are shown in gray. Notably, scaffold and synergy-based metrics (e.g., TScaffold, PhiID Syn, PED Syn, O-Syn) show stronger associations with language and fluid intelligence, while redundancy-based and triangle metrics are more strongly associated with attention and verbal episodic memory.

Computational complexity of HOI metrics

The computational cost of the HOI metrics varies significantly, impacting their scalability for whole-brain analyses. Table S1 and Figure S10 summarize the asymptotic time complexity and empirical runtimes (measured on a Ryzen Threadripper 3970 workstation using 16 cores for $T = 1200$, $N = 200$ regions), highlighting the dominant operations for each framework.

Pairwise baseline: functional connectivity. FC relies on Pearson correlation, requiring $\binom{N}{2}$ inner products over T samples. Highly optimized matrix multiplication routines (e.g. using BLAS) keep the runtime under 10 s, establishing the practical lower bound against which all HOI metrics are assessed

PhiID, bivariate. The Φ ID framework decomposes functional interactions at the pairwise level ($K = 2$), using a Gaussian estimator to compute the full set of redundancy–synergy atoms for every ordered pair of regions. It requires computing a lagged cross-covariance matrix in $O(N^2T)$ operations. Extracting the 4×4 joint distribution and computing atoms for all $O(N^2)$ pairs takes $O(1)$ per pair. The overall complexity mirrors FC, with a slightly higher runtime (~ 20 s) due to the atom decomposition overhead.

Static scaffold: single Rips filtration. The static homological scaffold is derived from the full-session correlation matrix (estimated once in $O(N^2T)$), which is converted to a distance matrix and passed to a Vietoris–Rips persistent homology algorithm (Ripserer) to extract H_1 (one-cycle) features. The filtration contains up to $m = O(N^3)$ simplices (i.e., nodes, edges, triangles). While theoretically bounded by $O(m^3)$, algorithm optimizations in persistent homology yield near-linear practical runtimes (~ 15 s). Because the filtration is built only once from time-averaged connectivity, it remains highly tractable.

O-information. For O-information (and related TC/DTC/S-information), a covariance matrix is computed once ($O(N^2T)$), followed by enumeration of all triplets ($O(N^3)$), where each triplet contribution requires only constant-time operations under the Gaussian estimator (i.e. computation of entropy terms). This yields intermediate runtimes (Table S1) and becomes increasingly bottlenecked by enumeration as N grows (Fig. S10).

PED. PED follows a fundamentally different strategy: rather than a Gaussian approximation, it relies on explicit binary discretization of the time series ($B = 2$ states, obtained by thresholding at zero after z -scoring). For each of the $\binom{N}{3} = O(N^3)$ triplets (i, j, k) , the joint binary distribution over $2^3 = 8$ states is built by a single pass through the T time points, yielding $O(T)$ operations per triplet. Because the joint state space is constant (2^3 , independent of T and N), the subsequent Partial Information Decomposition via the sx-lattice [8] is $O(1)$ per triplet. The overall complexity is therefore $O(TN^3)$, with possible parallelization/JIT acceleration. The wall-clock time is ~ 370 s at $N = 200$.

Dynamic TDA (TScaff) TScaff generalizes the static scaffold by computing persistent homology at every time point. Precomputing triplet co-fluctuations costs $O(TN^3)$, and extracting homology T times adds $T \cdot O(m_t^3)$. This compounding cost diverges rapidly for $N > 100$, taking ~ 2700 s.

Practical implications and scalability. Overall, our benchmarks reveal a clear stratification: pairwise/bivariate metrics and static topology are the most tractable; O-information is intermediate; PED and especially dynamic topology are the most expensive. These differences become decisive at higher interaction orders ($K \geq 4$), where candidate multiplets grow as $\binom{N}{K}$. Even for Gaussian formulations (where each K -tuple contribution remains $O(1)$ after covariance estimation), enumeration scales as $O(N^K)$ and quickly becomes limiting. Accordingly, higher-order analyses are most feasible with targeted strategies such as anatomical/network preselection, sparsification, or sampling of candidate multiplets (as discussed in the main text).




Figure S10: **Execution time of HOI metrics for a representative fMRI subject.** Execution time (in seconds) for computing pairwise and HOI metrics on a single subject ($T = 1200$ time points) as a function of the number of Regions of Interest (ROIs), ranging from $N = 10$ to $N = 200$. Benchmarks were performed on a Ryzen Threadripper 3970 workstation (using 16 cores). The inset (log-scale) highlights the orders of magnitude difference between the approaches. While PED and Dynamic TDA (Tsscaffold) both scale with the number of triplets ($O(TN^3)$), tscaffold (~ 2700 s at $N = 200$) is nearly an order of magnitude slower than PED (~ 370 s). This divergence reflects the dominant computational cost: PED relies on efficient binary joint-counting ($B = 2$), whereas tscaffold incurs an additional term $T \cdot O(m_t^3)$ representing the matrix-reduction step required to compute persistent homology at every time point (where m_t is the number of simplices in the dynamic filtration at time t). Static metrics (static scaffold, O-info) and pairwise metrics (FC, PhiID) remain computationally tractable (< 2 minutes) even at high parcellation densities ($N = 200$), scaling primarily with the initial covariance estimation ($O(N^2T)$) or static triplet enumeration. Error bars over 10 iterations are smaller than the marker size.

Metric	Time Complexity	Runtime (N=200)	Dominant Computation
Functional connectivity (FC)	$O(N^2T)$	< 10 s	Correlation estimation
PhiID (Bivariate)	$O(N^2T)$	~ 20 s	Lagged covariance + pairwise atoms
Static scaffold	$O(N^2T) + O(m^3)$	~ 15 s	Single filtration + H_1 extraction
O-information (Red/Syn)	$O(N^2T + N^3)$	~ 90 s	Covariance + triplet enumeration
PED (Red/Syn)	$O(TN^3)$	~ 370 s	Triplet-wise joint-counting (binary)
Dynamic TDA (TScaff)	$O(TN^3) + T \cdot O(m_t^3)$	~ 2700 s	Per-time-point filtration + homology

Table S1: **Computational complexity and empirical runtime of network metrics.** Asymptotic time complexity for N regions and T time points, with empirical runtime obtained with a good workstation (Ryzen Threadripper 3970 with 16 cores used), processing a reference subject with $T = 1200$ and $N = 200$ ROIs. PED and dynamic TDA both scale as $O(TN^3)$ due to triplet enumeration, but dynamic TDA incurs additional per-time-point topological computation with larger constant factors. Indeed, PED uses binary discretization ($B = 2$), so the per-triplet state space is constant (2^3). For topology-based methods, m denotes the number of simplices included in the static filtration (e.g., vertices, edges, and triangles), and m_t denotes the number of simplices included at time point t in the dynamic filtration. The $O(m^3)$ and $O(m_t^3)$ terms refer to a worst-case upper bound for the matrix-reduction step used to compute persistent homology.

The consistency of order 3 and order 4 O-information

In order to assess the generality of our findings with respect to the interaction order considered, we extended the O-information analysis to include quadruplets (order 4). Among the higher-order metrics considered in this work, the O-information is uniquely suited for this extension due to its relatively low computational cost and analytical formulation, which allows a straightforward generalization to larger group sizes. This makes it possible to evaluate whether the redundancy- and synergy-dominated structures identified at order 3 persist when considering interactions among four regions.

For order 4, we computed the O-information across all quadruplets of regions and decomposed it into redundancy and synergy components using the same definitions adopted for triplets. From these values, we derived both node-level and edge-level summaries. Node-level measures were defined as the average interaction strength $\langle s_i \rangle$ of a region across all quadruplets in which it participates. Edge-level measures were obtained by aggregating contributions from quadruplets containing each pair of regions, allowing direct comparison with the order 3 edge-level representation.

Fig. S11 displays the spatial organization of order 4 redundancy and synergy at the node level, as well as the corresponding edge-level matrices. Redundancy patterns at order 4 closely resemble those obtained at order 3, preserving the large-scale spatial structure observed in the main manuscript. Synergy exhibits more heterogeneous organization but maintains coherent system-level structure.

To quantify the agreement across interaction orders, we computed Pearson correlations between order 3 and order 4 estimates. Redundancy shows near-perfect correspondence at both the node ($r = 1.00$) and edge ($r = 0.99$) levels, indicating that redundancy-dominated organization is highly stable with respect to interaction order. Synergy remains strongly correlated across orders (node: $r = 0.83$; edge: $r = 0.79$), though with greater variability, suggesting that synergistic interactions are more sensitive to the inclusion of additional regions.

These results demonstrate that the principal redundancy structure identified in the main manuscript generalizes to higher interaction orders, while synergy reveals additional order-dependent modulation. Overall, the extension to order 4 confirms that the large-scale organizational principles captured by the O-information are not restricted to triplet interactions.




Figure S11: **Consistency of redundancy and synergy across interaction orders at node and edge levels.** The two top left panels show node-level maps of redundancy (left) and synergy (right) computed from order 4 interactions, displayed on the cortical surface (blue–red color scale indicates the ranking of regional mean interaction strength $\langle s_i \rangle$, computed as the sum across all quadruplets). The two top right panels show the corresponding edge-level redundancy and synergy matrices computed from order 4 interactions. Bottom panels compare order 4 versus order 3 estimates for redundancy and synergy at the node level (left) and edge level (right). Scatter plots report Pearson correlation coefficients (r), demonstrating near-perfect agreement for redundancy at both node ($r = 1.00$) and edge ($r = 0.99$) levels, and strong but reduced agreement for synergy (node: $r = 0.83$, edge: $r = 0.79$). These results indicate that redundancy patterns are highly stable across interaction orders, whereas synergy exhibits greater sensitivity to the inclusion of higher-order interactions, particularly at the edge level.

Persistent entropy allows to distinguish rest from task

To ensure that our conclusions regarding topological descriptors are as general as possible, we complemented the main analyses with persistent entropy, a global summary measure derived from persistent homology. As illustrated in Fig. S12a, resting-state and task fMRI time series (Schaefer–Tian atlas, $N = 116$) were used to compute correlation matrices, from which persistence diagrams were extracted across homological dimensions (H_0 , H_1 , H_2). Persistent entropy was then computed to quantify the overall topological complexity of these diagrams.

By including this additional topological metric, we assess whether behaviorally relevant differences between rest and task are detectable not only through interaction-based higher-order measures but also through global topological summaries of functional connectivity.

Figure S12b shows the distribution of persistent entropy values across subjects and conditions. Although entropy values occupy a relatively compact region of the feature space, systematic shifts are observable between resting-state and task conditions across homological dimensions. This indicates that task engagement induces measurable changes in the global topological organization of functional connectivity.

To assess the discriminative power of these features, we performed task decoding analyses using persistent entropy values as input features. Classification accuracy was evaluated using a support vector machine (SVM) classifier, as in the main task decoding analyses. Both binary discrimination (rest versus task) and multiclass discrimination across all task conditions were considered. As shown in Fig. S12c, persistent entropy achieves high accuracy in distinguishing rest from task, indicating that global topological summaries carry discriminative information about cognitive state. These results demonstrate that topological summaries capture meaningful differences in large-scale brain dynamics associated with cognitive state.

Together, these findings indicate that persistent entropy provides a compact yet informative descriptor of functional brain organization, sensitive to task-related modulations and suitable for decoding analyses.

Test with additional pairwise and higher-order metrics

To provide a broader methodological coverage of interaction measures and to further characterize the relationship between pairwise and higher-order descriptors, we extended the comparative framework by incorporating additional pairwise and higher-order metrics. The aim of this extension is twofold: first, to provide a more systematic investigation of higher-order metrics, incorporating a recently proposed regression-based method that is conceptually distinct from the information-theoretic and topological frameworks considered in the main text; and second, to assess the robustness of the observed clustering structure and fingerprinting performance upon inclusion of additional metrics.

Additional pairwise metrics.

In addition to functional connectivity (FC), we considered mutual information (MI), transfer entropy (TE), and precision (inverse covariance) as complementary pairwise descriptors.

MI and TE were computed using the HOI toolbox, which provides standardized implementations for estimating information-theoretic dependencies from multivariate time series. For both MI and TE, we used a Gaussian-copula estimator; for TE, we used a lag of 1 time step. MI quantifies zero-lag statistical dependence between pairs of regional BOLD signals, capturing both linear and nonlinear associations. TE, in contrast, measures directed and time-lagged information transfer, thereby incorporating temporal asymmetry into the pairwise representation. These measures allow the characterization of statistical dependencies beyond linear correlation while remaining restricted to pairwise interactions.

Precision (i.e., the inverse covariance matrix) was estimated using the PySPI toolbox with the empirical covariance estimator (dense precision matrix)[9]. Its inclusion provides a reference pairwise model




Figure S12: **Persistent entropy of brain dynamics across rest and task conditions.** (a) Analysis pipeline. Resting-state and task fMRI data from the Human Connectome Project (HCP) are parcellated using the Schaefer–Tian atlas ($N = 116$ regions), yielding regional time series. Persistent homology is computed on the correlation matrix obtained from the resulting functional time series, producing persistence diagrams whose topological complexity is summarized using persistent entropy. (b) Persistent entropy values across homological dimensions (H_0 , H_1 , H_2) for all subjects and conditions. Each point corresponds to a subject–task instance, colored by task. While persistent entropy values occupy a compact manifold, differences between rest and task are evident across dimensions. (c) Task decoding performance using persistent entropy features. Classification accuracy is shown for discriminating rest versus task conditions and for multiclass decoding across all tasks. Boxplots indicate the median and interquartile range across bootstrap resamples, with individual points representing bootstrap-level accuracies. Persistent entropy reliably distinguishes rest from task states and supports above-chance decoding across multiple task conditions, indicating that topological summaries capture behaviorally relevant differences in large-scale brain dynamics.

that accounts for conditional relationships and has been shown by recent literature to be the pairwise metric with the highest score in brain fingerprinting. Given that higher-order measures exhibited stronger performance in brain fingerprinting compared to pairwise approaches, and that precision has been identified in previous work as the pairwise metric achieving the highest fingerprinting accuracy, we included precision in the present analysis to provide a rigorous and state-of-the-art benchmark for comparison.

Regression-based higher-order inference (SINDy).

To complement the information-theoretic and topological approaches, we incorporated a regression-based higher-order inference method grounded in the Sparse Identification of Nonlinear Dynamics (SINDy) framework, following its recent formulation for hypergraph inference [10]. In line with the original paper, we estimated node-wise derivatives from the time series via finite differences, expanded the dynamics using a Taylor basis up to degree (2), and used sequential thresholded least-squares regression to infer sparse coefficients, retaining nonzero terms as inferred hyperedge weights. We also followed the paper’s pre-filtering strategy, restricting candidate terms using only the strongest pairwise correlations (top 10% of ROI pairs).

In this framework, higher-order hyperedges of order k are inferred by modeling the joint contribution of $k - 1$ regional time series to the future activity of the target region using sparse regression. The estimated regression coefficients are interpreted as weights of k -order hyperedges. Unlike entropy-based or topological summaries, this approach identifies higher-order dependencies through a dynamical systems perspective, without relying on information decompositions or persistence diagrams. The same projection procedures described in the original study were adopted to ensure methodological consistency.

Brain fingerprinting performance.

Fingerprinting analyses were repeated using the expanded set of metrics. At the whole-brain level, higher-order descriptors maintain the strongest identifiability performance across two evaluation criteria. Scaffold-based measures exhibit the largest identifiability effect size (Cohen’s d), while Φ ID-red achieves the highest identification rate, jointly with MI. At the level of resting-state network pairs, higher-order metrics outperform pairwise approaches in 80 out of 108 comparisons. This consistency across global and network-resolved analyses indicates that the advantage of higher-order representations in capturing subject-specific structure is robust to the inclusion of additional pairwise metrics.

Similarity structure of metrics.

The inclusion of these additional measures allows a re-evaluation of the edge-level similarity structure across all descriptors. Spearman correlations between vectorized connectivity representations were computed to quantify pairwise similarity between metrics. The extended similarity matrix preserves the previously observed clustering structure. Redundancy-based measures exhibit high similarity with MI and SINDy, suggesting that these approaches capture related aspects of shared information structure. TE shows closer alignment with redundancy-dominated metrics, whereas precision displays partial similarity with synergy-related descriptors. Topological and scaffold-based measures remain comparatively distinct.

Hierarchical clustering of the full similarity matrix confirms the presence of three principal families of descriptors: redundancy-dominated, synergy-dominated, and topological/scaffold-based metrics. The integration of the newly considered measures within these pre-existing clusters indicates that the organizational structure of higher-order descriptors is stable across methodological extensions.




Figure S13: Extended comparison of connectivity metrics for brain fingerprinting. (a) Whole-brain fingerprinting performance evaluated across an expanded set of pairwise and higher-order metrics. Identification rate (left), individual identifiability score I_{diff} (middle), and effect size (Cohen's d , right) are reported to provide complementary assessments of brain fingerprinting. While precision, mutual information, and transfer entropy improved the performance of pairwise approaches, higher-order descriptors, particularly scaffold-based and redundancy-related measures, consistently achieve greater fingerprinting scores across metrics. Box plots show the median, interquartile range (IQR), and $1.5 \times \text{IQR}$ whiskers across 100 bootstrap iterations (80 subjects selected out of 100). (b) Network-level fingerprinting performance computed within and between the seven Yeo resting-state networks. The best-performing metric varies across systems, indicating that subject-specific information is heterogeneously distributed across functional networks.




Figure S14: **Edge-level similarity and hierarchical clustering, including additional higher-order and pairwise connectivity metrics.** (a) Edge-level similarity matrix across all evaluated metrics, quantified using Spearman correlation between vectorized connectivity representations. Each entry reports the similarity between two metrics computed as the Spearman r across all edges. Redundancy-based measures (Red, O-Red, PED Red, Φ ID Red), together with mutual information (MI) and SINDy, show very high mutual similarity, forming a tightly correlated group. Transfer entropy (TE) and precision (prec) exhibit slightly lower correlations, with TE aligning more closely with redundancy-based metrics and precision showing greater similarity to synergy-related descriptors. (b) Hierarchical clustering dendrogram derived from the edge-level similarity matrix. The clustering confirms the presence of three main families of descriptors: redundancy-dominated metrics (including MI, SINDy, and TE), synergy-dominated metrics, and topological/scaffold-based measures. Importantly, the newly included metrics integrate consistently within the pre-existing clusters, reinforcing the robustness of the previously identified organizational structure.

Hypergraph metrics to compute brain regions centrality

In the main manuscript, node importance within higher-order representations of brain activity was primarily quantified using strength centrality, defined as the sum (or average) of the weights of all hyperedges incident to a given region. While this measure provides an intuitive and computationally efficient characterization of nodal involvement in higher-order interactions, it does not explicitly account for recursive effects arising from the higher-order topology of the hypergraph.

To further characterize the role of brain regions within the inferred hypergraphs, we therefore considered centrality measures that have been recently generalized from pairwise networks to hypergraphs. In particular, we computed H-eigenvector centrality [11] and Katz (α -) centrality, as implemented in the `xgi` Python toolbox [12], which provides a comprehensive framework for the construction and analysis of higher-order networks, including hypergraph centrality algorithms¹.

H-eigenvector centrality extends classical eigenvector centrality to uniform hypergraphs by defining node importance through a nonlinear eigenvalue problem on the adjacency tensor of the hypergraph [11]. Katz centrality (also referred to as α -centrality in the hypergraph setting) generalizes walk-based centrality measures to higher-order structures by accounting for contributions from walks of all lengths, attenuated by a damping parameter.

For each higher-order construction (including O-Red, O-Syn, PED Red, PED Syn, and Triangles, inferred hypergraphs [10]), we computed nodal strength, H-eigenvector centrality, and Katz centrality. Before computing these centrality measures, we applied a binarization threshold at the 60th percentile of the hyperedge weight distribution (we also reproduced these findings with other thresholds, namely 70th, 80th, and 90th percentiles). Pearson correlations between these centrality measures were evaluated across brain regions to assess their consistency and relative behavior. Across all higher-order interaction models, hypergraph centralities exhibited consistently high correlations with nodal strength (typically $r \approx 0.8 - 1.0$), as well as strong agreement with each other. This indicates that, in the present setting, strength centrality captures a substantial portion of the structural information encoded by more nuanced hypergraph centrality formulations. While future work may explore a more detailed comparison of alternative hypergraph centrality formulations, the present analysis indicates that nodal strength, as computed in the main manuscript, is broadly aligned with more sophisticated centrality measures. This suggests that the reported results are likely to generalize across different definitions of hypergraph centrality.

¹<https://xgi.readthedocs.io/en/latest/api/algorithms/xgi.algorithms.centrality.html>





Figure S15: **Comparison of hypergraph centrality measures across higher-order constructions.** Pearson correlation matrices between nodal strength, H-eigenvector centrality, and Katz (α -) centrality computed on hypergraphs inferred using different higher-order metrics (Red, PED red, PED syn, Syn, triangles), together with their average across constructions. Each panel reports the pairwise correlation between centrality measures across brain regions. Across all higher-order definitions, hypergraph centralities show consistently high correlations with nodal strength and with each other ($r \approx 0.8$ -1.0), indicating strong agreement in the ranking of regional importance. These results demonstrate that the strength-based centrality used in the main manuscript captures a substantial portion of the structural information encoded by more sophisticated hypergraph centrality formulations, while confirming the robustness of node-level characterizations across different higher-order interaction models.



Supplementary References

- [1] Lianne H. Scholtens, Siemon C. de Lange, and Martijn P. van den Heuvel. Simple brain plot, “Zenodo”, <https://doi.org/10.5281/zenodo.5346593>, 2021.
- [2] 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, September 2018.
- [3] Ye Tian, Daniel S. Margulies, Michael Breakspear, and Andrew Zalesky. Topographic organization of the human subcortex unveiled with functional connectivity gradients. *Nature Neuroscience*, 23(11):1421–1432, November 2020.
- [4] Michael D Greicius, Ben Krasnow, Allan L Reiss, and Vinod Menon. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proceedings of the national academy of sciences*, 100(1):253–258, 2003.
- [5] Andrea I. Luppi, Pedro A.M. Mediano, Fernando E. Rosas, Negin Holland, Tim D. Fryer, John T. O’Brien, James B. Rowe, David K. Menon, Daniel Bor, and Emmanuel A. Stamatakis. A synergistic core for human brain evolution and cognition. *Nature Neuroscience*, 25:771–782, 6 2022.
- [6] Valerie J Sydnor, Bart Larsen, Danielle S Bassett, Aaron Alexander-Bloch, Damien A Fair, Conor Liston, Allyson P Mackey, Michael P Milham, Adam Pines, David R Roalf, et al. Neurodevelopment of the association cortices: Patterns, mechanisms, and implications for psychopathology. *Neuron*, 109(18):2820–2846, 2021.
- [7] Bertha Vázquez-Rodríguez, Laura E. Suárez, Ross D. Markello, Golia Shafiei, Casey Paquola, Patric Hagmann, Martijn P. van den Heuvel, Boris C. Bernhardt, R. Nathan Spreng, and Bratislav Misic. Gradients of structure–function tethering across neocortex. *Proceedings of the National Academy of Sciences*, 116(42):21219–21227, October 2019.
- [8] Abdullah Makkeh, Aaron J. Gutknecht, and Michael Wibral. Introducing a differentiable measure of pointwise shared information. *Phys. Rev. E*, 103:032149, Mar 2021.
- [9] Oliver M. Cliff, Annie G. Bryant, Joseph T. Lizier, Naotsugu Tsuchiya, and Ben D. Fulcher. Unifying pairwise interactions in complex dynamics. *Nature Computational Science*, 3(10):883–893, Oct 2023.
- [10] Robin Delabays, Giulia De Pasquale, Florian Dörfler, and Yuanzhao Zhang. Hypergraph reconstruction from dynamics. *Nature Communications*, 16(1):2691, March 2025.
- [11] Austin R Benson. Three hypergraph eigenvector centralities. *SIAM Journal on Mathematics of Data Science*, 1(2):293–312, 2019.
- [12] Nicholas W Landry, Maxime Lucas, Iacopo Iacopini, Giovanni Petri, Alice Schwarze, Alice Patania, and Leo Torres. Xgi: A python package for higher-order interaction networks. *Journal of Open Source Software*, 8(85):5162, 2023.