RESEARCH ARTICLE

Interacting brains revisited: A cross-brain network neuroscience perspective

Christian Gerloff[1,2,3] (iD) | Kerstin Konrad[1,2] | Danilo Bzdok[4,5] | Christina Büsing[3] | Vanessa Reindl[1,2,6]

[1]JARA-Brain Institute II, Molecular Neuroscience and Neuroimaging, RWTH Aachen & Research Centre Juelich, Aachen, Germany

[2]Child Neuropsychology Section, Department of Child and Adolescent Psychiatry, Psychosomatics and Psychotherapy, Medical Faculty, RWTH Aachen University, Aachen, Germany

[3]Chair II of Mathematics, Faculty of Mathematics, Computer Science and Natural Sciences, RWTH Aachen University, Aachen, Germany

[4]Department of Biomedical Engineering, McConnell Brain Imaging Centre, Montreal Neurological Institute, Faculty of Medicine, McGill University, Montreal, Canada

[5]Mila – Quebec Artificial Intelligence Institute, Montreal, Canada

[6]Psychology, School of Social Sciences, Nanyang Technological University, Singapore, Singapore

Correspondence
Christian Gerloff, JARA-Brain Institute II, Molecular Neuroscience and Neuroimaging, RWTH Aachen & Research Centre, Juelich, Germany.
Email: christian.gerloff@rwth-aachen.de

Abstract

Elucidating the neural basis of social behavior is a long-standing challenge in neuroscience. Such endeavors are driven by attempts to extend the isolated perspective on the human brain by considering interacting persons' brain activities, but a theoretical and computational framework for this purpose is still in its infancy. Here, we posit a comprehensive framework based on bipartite graphs for interbrain networks and address whether they provide meaningful insights into the neural underpinnings of social interactions. First, we show that the nodal density of such graphs exhibits nonrandom properties. While the current hyperscanning analyses mostly rely on global metrics, we encode the regions' roles via matrix decomposition to obtain an interpretable network representation yielding both global and local insights. With Bayesian modeling, we reveal how synchrony patterns seeded in specific brain regions contribute to global effects. Beyond inferential inquiries, we demonstrate that graph representations can be used to predict individual social characteristics, outperforming functional connectivity estimators for this purpose. In the future, this may provide a means of characterizing individual variations in social behavior or identifying biomarkers for social interaction and disorders.

KEYWORDS
Bayesian modeling, brain-to-brain synchrony, fNIRS, graph machine learning, hyperscanning, interbrain networks, network embeddings

1 | INTRODUCTION

Network models have provided new insights into the neural basis of social behavior. Traditionally, neuroscience studies have examined functional connectivity within the brain of a single subject using standardized tasks. These studies have yielded mounting evidence that the integrative processes and dynamic interplay among different brain regions and systems underpin many of our cognitive processes (van

den Heuvel & Sporns, 2013). While these studies have focused on single brains, each person is embedded in a larger network of social partners and continuously affects and is affected by others (Falk & Bassett, 2017; Hari et al., 2015). Thus, to reach a deeper comprehension of the neural underpinnings of social interactions, it may not be sufficient to study individual brains in isolation (Hari et al., 2015). Instead, the continuous reciprocal exchange between interacting brains should be the target of investigation.

Studies that record the brain activities of two or more persons concurrently, a mode of investigation referred to as hyperscanning, indicate that the brain activities of interacting persons become synchronized, that is, show statistical dependencies (interpersonal neural synchrony [INS]) (Babiloni & Astolfi, 2014). However, many of these studies have analyzed INS only between region pairs of the two participants, and often in homologous brain regions, thereby neglecting the fact that each brain is organized and operates as a complex system (e.g., Cui et al., 2012; Djalovski et al., 2020; Reindl et al., 2018). Hence, for research questions that do not explicitly target a priory specified brain regions, such pairwise comparisons of region pairs may not be sufficient to fully capture INS, given the large number of dynamic dependencies between brain regions as well as individual variability in functional networks. While the first few, mostly electroencephalography (EEG)-based, hyperscanning studies have adopted a graph analytic approach to study INS (e.g., Ciaramidaro et al., 2018; Santamaria et al., 2020), graph analysis for hyperscanning data is still in its infancy (Toppi et al., 2015). Here, we posit a comprehensive analytical framework for inference and prediction based on bipartite interbrain graphs. Specifically, we extend current analyses by (i) considering both global and nodal topological properties, (ii) comparing the properties of interbrain graphs to those of random graphs, and (iii) exploring whether graph representations can be used to infer relationships and predict social characteristics.

Conceptually, synchrony can be understood as the coordinated or aligned temporal relationship between events or time series. To quantify the statistical dependencies between neural signals, a variety of methods exist, each under its own premises (see Bastos & Schoffelen, 2016). Based on these estimators, a bipartite graph can be constructed from two sets of nodes representing each person's brain regions and edges encoding their statistical dependencies. Importantly, such a bipartite graph enables full preservation of spatial precision in two nonoverlapping region-of-interest (ROI) sets, which sets it apart from traditional seed-based connectivity mapping, in which the seed ROI must be averaged, causing its spatial information to be lost. From such a bipartite graph, both global and nodal properties can be derived. Previous graph-based hyperscanning studies have often focused on global metrics (e.g., Ciaramidaro et al., 2018; Santamaria et al., 2020). A global metric, such as global efficiency, is a single metric (scalar) that aggregates a specific graph topological property across the entire graph (Latora & Marchiori, 2001). Hence, global metrics cannot provide insights into the local differences in a graph's topology and may thereby overlook smaller and/or more localized effects. Thus, understanding how connectivity varies across nodes is an important step in network analysis (Fornito et al., 2016). One of the most fundamental nodal metrics, to which most other network metrics are ultimately linked, is the nodal degree, that is, the number of edges per node, or the nodal density, that is, the degree normalized with respect to the number of possible edges (Bullmore & Sporns, 2009).

1.1 | Local topology of real-world networks

Nodal metrics not only provide deeper insights into network organization but were also used in early network neuroscience studies to assess the general properties and feasibility of graph formulations (e.g., Achard et al., 2006). In an intrabrain network, as in many other large, real-world networks, not every node is equally important to the network's topology. While most nodes have very few connections (a low degree), some nodes have many connections (a high degree), facilitating integration across the network (Barabási & Albert, 1999; Fornito et al., 2016). Consequently, the degree distribution shows heavy-tailed characteristics and differs from the degree distribution of a random graph with uniform edge probabilities, such as an Erdős-Rényi graph (Erdös & Rényi, 1959). However, these fundamental topological properties have so far not been examined for interbrain networks, although they may provide first evidence for the validity of a network-based perspective on INS. Thus, in this study, we investigate whether interbrain networks show similar heavy-tailed properties, indicating the presence of seed regions, and importantly, whether their nodal density distributions differ from those of random graphs.

1.2 | Bayesian inference based on interbrain network representations

Nodal metrics offer an in-depth view of a graph's topology. However, they come at the cost of high computational complexity and multiple comparison issues. Hence, an approach that can yield localized, interpretable insights into interbrain networks while avoiding multiple comparison issues is currently lacking. To address this problem, we use an unsupervised machine learning technique, the nonnegative matrix factorization (NMF; Lee & Seung, 1999), to decompose an interbrain network into several collections of nodes. This allows us to encode the nodal topologies in a low-dimensional vector while preserving the contributions of the individual nodes to each component. The coefficients of these components are then analyzed within a multivariate Bayesian hierarchical model (BHM), thereby integrating all nodal effects into a single model and avoiding multiple comparisons. Furthermore, Bayesian analysis goes beyond the point estimates of a frequentist approach by allowing a probability statement to be derived for each parameter at both the individual and population levels (see also Kruschke, 2014).

1.3 | Interbrain network-based prediction

The Bayesian model explicitly follows the explanatory goal of inference for building and testing neuroscientific claims (see also

Shmueli, 2010). While many previous hyperscanning studies have used inferential modeling, an important long-term goal may be to predict behavioral responses, clinical phenotypes, or treatment outcomes. In addition, predictive modeling provides a straightforward way to quantify the performance of our NMF-based approach by comparing its ability to predict known properties of the data to the corresponding capabilities of other approaches. This enables us to address whether interbrain networks provide an informative representation that goes beyond functional connectivity estimators.

Prediction based on vectorized connectivity estimators (see also "connectivity fingerprinting" Finn et al., 2015; Passingham et al., 2002) captures only pairwise dependencies while neglecting the multiple dependencies among regions of the same network. In contrast, predictive models for classifying graph-structured data represent a new direction in neuroscience (Rosenthal et al., 2018). The underling computational problem of such graph comparison is well known in graph theory; formally, under the premise to determine whether two graphs are topologically equivalent, termed graph isomorphism problem (Fortin, 1996). Going beyond first attempts, such as the pairwise Weisfeiler-Lehman test, graph representation learning (GRL) is gaining enormous traction (Goyal & Ferrara, 2018; Wu et al., 2021). The aim of GRL is to learn low-dimensional continuous vector representations for graph-structured data, called embeddings (Bronstein et al., 2017; Chami et al., 2022). These methods involve embedding nodes or graphs that are similar with respect to a certain association into a graph space and then transforming them into Euclidean space (Chami et al., 2022). Since our NMF-based approach preserves the topological properties of a graph, it can be understood as a "structural embedding" of a graph. However, since our approach is based on topological properties, other network embeddings based on other lower- or higher-order proximities could potentially yield superior network representations. Hence, in a rigorous testing regime, we compare the performance of our network representation with state-of-the-art network embeddings (Graph2Vec: Narayanan et al., 2017; GL2Vec: Chen & Koga, 2019; GeoScattering: Gao et al., 2019; LDP: Cai & Wang, 2018) and predictions based on functional connectivity estimators.

We demonstrate the proposed analytical framework and evaluate its utility using a functional near-infrared spectroscopy (fNIRS) hyperscanning data set with a well-established hyperscanning design (for more details, see Reindl et al., 2018). fNIRS measures the hemodynamic response in cortical brain regions, usually with a higher spatial resolution than EEG, and is based on changes in the concentrations of both oxy- and deoxyhemoglobin (HbO and HbR). The latter property is an advantage of fNIRS because it allows us to conduct analyses based on the HbO signals, as is most commonly done (e.g., Cui et al., 2012; Reindl et al., 2018), while using the HbR signals as a hold-out data set for the prediction task to validate the findings.

In the current study, female children and adolescents performed both a cooperative and a competitive computer task (with two task blocks each) with their mother and with an adult stranger, while their prefrontal brain activities were recorded concurrently. The two task conditions were compared to a baseline condition in which both participants watched a relaxing video together. The mechanisms that drive synchrony may vary between experimental tasks and are thought to span different behavioral and biological systems (Feldman, 2012; Hamilton, 2021; Reindl et al., 2022). For the given experimental design, it can be hypothesized that, on the one hand, parent–child attachment is associated with processes of increased bio-behavioral synchrony (Feldman, 2012) and, on the other hand, that the mutual engagement in a given task induces similar cognitive and emotional processes in the dyad which are time-locked to each other (see Reindl et al., 2022 for more details). Given the adolescent sample, an age in which children strive for more autonomy from their parents and orient more toward peers, social cooperation, and competition with parents, peers, and strangers may play an important role for their socioemotional development (see Kruppa et al., 2021; Nelson et al., 2005). Previous research based on the same experimental task enables a plausibility check of the results: based on previous findings, stronger INS is expected for mother–child dyads than for stranger–child dyads (Kruppa et al., 2021; Reindl et al., 2018) and for the two task conditions compared to the baseline condition in the dorsolateral and frontopolar brain regions (Kruppa et al., 2021).

2 | MATERIALS AND METHODS

2.1 | Empirical data

2.1.1 | Participants

Data were collected as part of a larger project (see Reindl et al., 2022), in which the same fNIRS data were analyzed to investigate whether neural synchrony goes beyond synchrony in the autonomic nervous system and motor behavior using multimodal models. Participants were recruited via previous studies, postings on the intranet of the University Hospital RWTH Aachen and flyers, and they were screened for severe cardiac, neurological or psychiatric conditions prior to study inclusion. The final sample consisted of 34 children (all female) who participated in the study with their mothers (mother–child dyads) and a previously unacquainted female adult (stranger–child dyads). A total of 29 strangers were included, some of which participated twice ($n = 1$) and some three times ($n = 2$), thus, comprising 34 mother–child and 34 stranger–child dyads. The average age of the children was 14.26 years (SD = 2.21 years, range: 10–18 years), that of the mothers was 45.32 years (SD = 4.95 years, range: 37–56 years), and that of the strangers was 23.07 years (SD = 2.09 years, range: 19–29 years). The strangers were significantly younger than the mothers (t(61) = −22.532, $p < .001$). All adults gave written informed consent for their own study participation and, in the case of the mothers, for the participation of their children. The children gave written informed assent (aged 10–17 years) or informed consent (aged 18 years). Ethics approval was obtained from the Ethics Committee of the Medical Faculty, University Hospital RWTH Aachen (EK 151/18).

2.1.2 | Data acquisition

fNIRS signals were recorded from both members of each dyad concurrently using a single fNIRS device (ETG-4000, Hitachi Medical Corporation, Japan; sampling rate: 10 Hz). Fifteen probes, eight sources and seven detectors, were arranged in grids consisting of three rows and five probe columns, in which the source-detector separation was fixed at 3 cm and the sources and detectors were arranged in alternating order, resulting in 22 measurement channels (CHs). Probe grids were mounted on modified EEG caps (Easycap GmbH, Germany) and placed symmetrically on the participant's foreheads, with the middle probe of the lowest probe row placed on the Fpz point in the 10–20 system and the middle probe column aligned along the sagittal reference curve. This probe setup, depicted in Figure 1, covered parts of Brodmann areas (BAs) 8, 9, 10, and 46. The most likely Montreal Neurological Institute (MNI) coordinates of the probes and channels as well as their underlying brain regions were estimated via the virtual registration method (Singh et al., 2005; Tsuzuki et al., 2007), using the Talairach Daemon (Lancaster et al., 2000). The most likely MNI coordinates of the fNIRS channels are provided in the corresponding repository (Gerloff, 2022). For more details on the data acquisition, see Reindl et al. (2019).

2.1.3 | Tasks and procedure

The participants in each dyad were seated next to each other in front of a single computer screen. A towel was placed over their hands to reduce their ability to view each other's movements. Furthermore, the participants rested their chins on chin rests to reduce head movements and improve the fNIRS data quality. The experiment consisted of three experimental conditions: a noninteractive baseline, a cooperative task, and a competitive task. Each child completed the three conditions either first with the mother ($N = 17$, 50%) and then, after a short break, with the stranger or the other way around ($N = 17$, 50%). The three conditions were presented in the same order for both dyads of which each child was a member, always starting with the baseline condition, followed either by the cooperative ($N = 16$, 47%) or the competitive ($N = 18$, 53%) task. A more detailed task description can be found in Text S1.

Baseline

A 3-minute excerpt from a relaxing aquatic video (Coral Sea Dreaming, Small World Music Inc.) was presented. This served as a control condition, in which both participants perceived the same stimuli and were instructed not to interact with each other.

Cooperation and competition

The cooperative and competitive computer games were adapted versions of the tasks of Cui et al. (2012), modified for use in children (Kruppa et al., 2021; Reindl et al., 2018). Each player manipulated a cartoon figure of a dolphin by pressing a computer key to either catch a ball together (cooperation) or catch the ball faster than the other

player (competition). The cooperative and competitive games were both organized into two task blocks of 20 trials each, interspersed with 30 s rest blocks: rest1, task1, rest2, task2, rest3.

2.2 | fNIRS preprocessing

First, the raw intensity data were converted into optical density data. Second, motion artifacts were detected using a moving standard deviation and reduced through cubic spline interpolation of the affected part of the signal (Scholkmann et al., 2010). Third, the optical density data were converted into HbO and HbR concentration changes using the modified Beer–Lambert law, in which the differential pathlength factor was individually estimated based on the wavelength and the participant's age (Scholkmann & Wolf, 2013). Finally, the HbO/HbR time series were detrended using a discrete cosine transform set. Noisy fNIRS channels were identified based on a semiautomated procedure, using several objective criteria (coefficient of variation, "flat line" detection, anticorrelation of HbO/HbR) in combination with visual inspection of the HbO/HbR signals and their scalogram. Based on this procedure, 7.84% of signal pairs were excluded from further analyses. If more than 25% of a participant's channels under a specific experimental condition were classified as noisy, the complete fNIRS measurement was excluded (see Text S2).

2.3 | Graphs

2.3.1 | Connectivity estimator

To formulate the interbrain networks, first the statistical dependencies between the dyad's brain signals have to be quantified. To this end, the bivariate wavelet coherence (WCO; Torrence & Compo, 1998) was calculated via the continuous wavelet transformation, which is a widely applied functional connectivity estimator that localizes signal dependencies in the time-frequency space and is appropriate for nonstationary time-series, such as fNIRS signals. The WCO belongs to the family of undirected connectivity estimators, which, unlike directed estimators such as Granger causality, do not allow to deduce the directionality of the relationship between neural time series (see Bastos & Schoffelen, 2016). Given its well-established mathematical foundations, its fit to signal characteristics and interpretability, we decided to exemplify this framework using the widely applied WCO. For each WCO matrix, the coefficients were aggregated into a scalar representing the connectivity estimator. To increase the robustness of the estimator, we applied padding to reduce boarder dispersions and considered only salient WCO coefficients that were higher than a cut-off value, determined on the basis of surrogate time-series, and lay within the cone-of-influence (for more information on the calculation of the WCO, the surrogates and the cut-off value see Text S3). Finally, we calculated the percentage of salient values within a task-related frequency band between 0.5 and 0.08 Hz (period length: 2.02–12.80 s) for the baseline condition

as well as each of the two cooperative/competitive task blocks. The task-related frequency band was chosen based on previous studies (Kruppa et al., 2021; Reindl et al., 2018) and included the trial duration (~7 s for cooperation, ~6 s for competition).

2.3.2 | Graph construction

To study interbrain networks and their global properties, we first formulated a complete bipartite graph for each dyad as depicted in
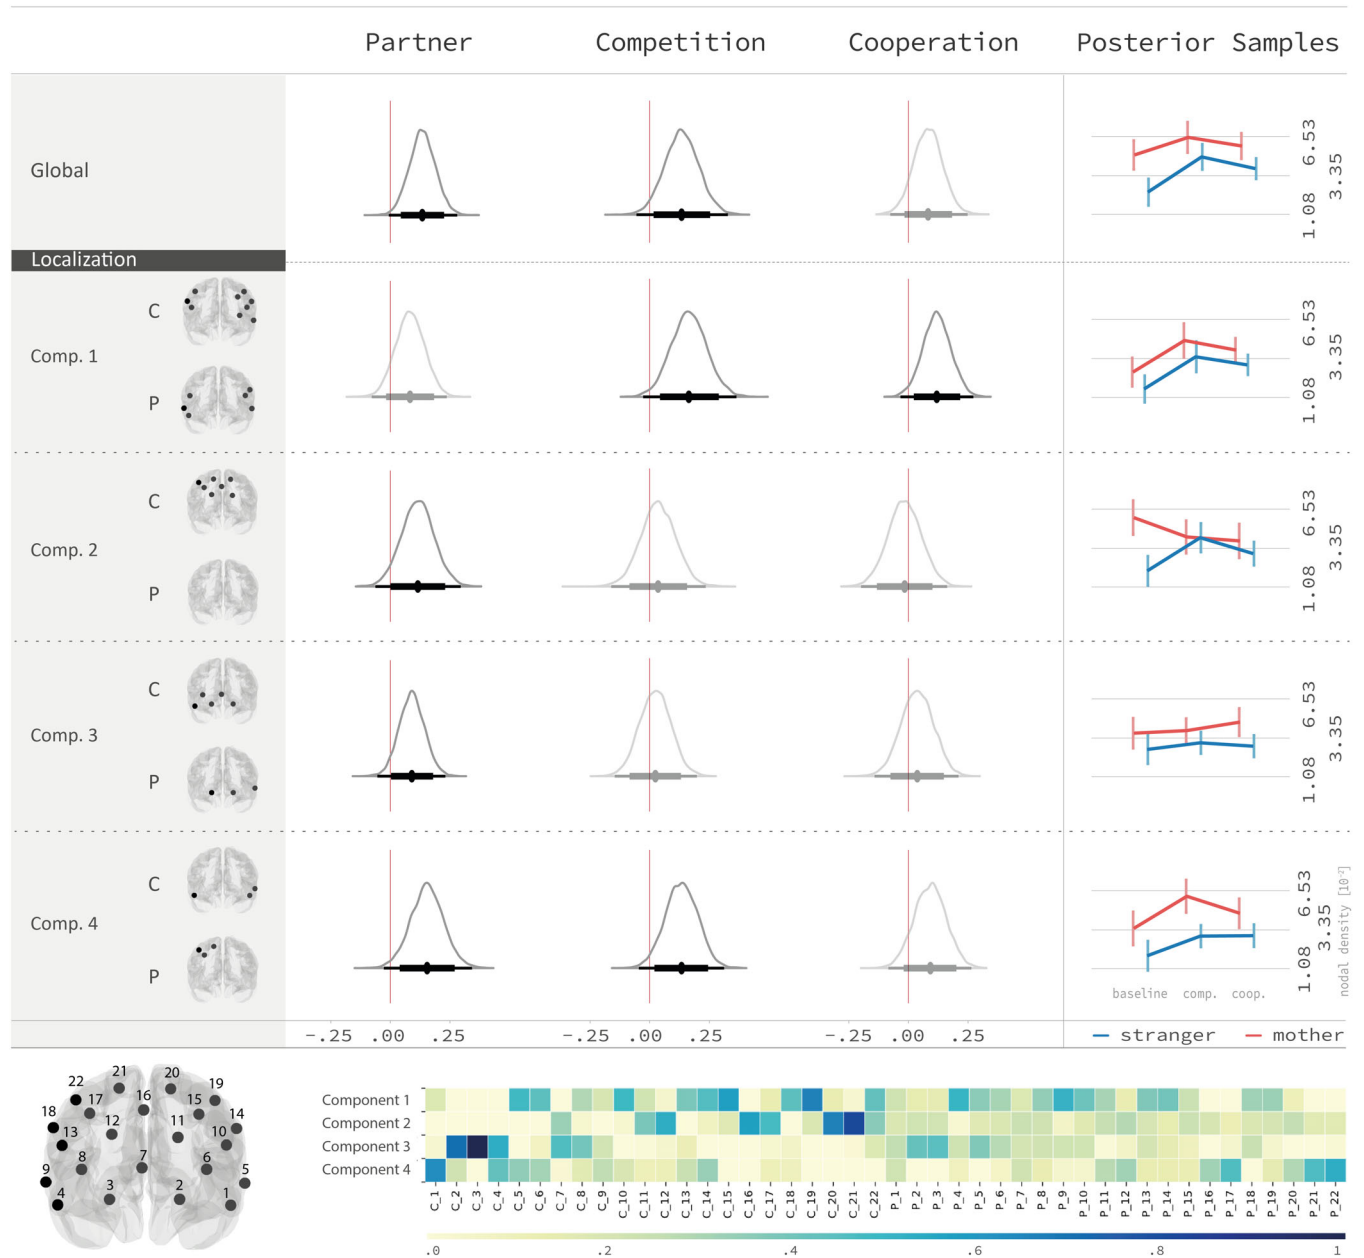
FIGURE 1 **Global effects of interbrain networks can be decomposed into nodal effects in specific seed regions.** This figure shows the global and nodal density effects of stranger-child interaction vs. mother–child interaction, baseline vs. competition and baseline vs. cooperation. The nodal densities were encoded by the coefficients of the four NMF components. Left: The marginal posterior distributions of the population effects are plotted along with their means, 90% CIs (thick black lines) and 99% CIs (thin gray lines). The width of the CI represents the uncertainty associated with the estimated parameters. Parameters for which the 90% CI does not include zero (red line) were interpreted as evidence for an effect. The global results showed a strong partner effect, with increased synchrony for mother–child dyads compared to stranger-child dyads, as well as a competition effect and some evidence for a cooperation effect, with higher synchrony for the two task conditions compared to the baseline condition. The nodal results confirmed these global results but provided increased topological detail. Specifically, the partner effect was rather widespread (in components 3 and 4 as well as for the baseline condition, in component 2), while the competition and cooperation effects were mostly localized to left and right lateralized prefrontal brain regions (components 1 and 4). Right: Five catplots show how the posterior global and nodal density values vary between mother–child and stranger-child dyads in the baseline, cooperation and competition conditions. Bottom: A heatmap visualizes the basis matrix resulting from NMF, showing the contribution of each fNIRS channel of the child (C) and the adult partner (P) (x-axis) to the corresponding component (y-axis). The fNIRS channels of the child and the adult partner that contribute most to each of the components in terms of their nodal densities, with weights above the 80th percentile (min = 0, max = 1), are visualized on the brain models

Figure 2. A bipartite graph $G = (V_1 \cup V_2, E)$ was constructed from the connectivity between the brain region of participant 1 and the brain regions of participant 2. The nodes comprising the two disjoint sets V_1 and V_2, also termed partite sets, represented the spatially localized sources for participant 1 and participant 2, respectively. $E \subseteq V_1 \times V_2$ denotes the edges with the corresponding weights W, defined by the connectivity estimator (Section 2.3.1). For each dyad and task, a complete bipartite graph $G_{complete}$ was generated with $|V_1| = |V_2| = 22$ nodes. Finally, edges from noisy channels were excluded. Thus, after exclusion of noisy data sets, a total of $n = 330$ graphs based on HbO signals and $n = 330$ graphs based on HbR signals were constructed. For further implementation details see Text S4.

2.3.3 | Graph reduction

One well-known issue in graph analysis is that complete graphs are sensitive to errors in the derived edges, leading to spurious connections and high false positive rates (van den Heuvel et al., 2017). Hence, complete graphs are often reduced via proportional thresholding; however, this has been shown to preserve some spurious connections, particularly in data sets with a low overall functional connectivity (van den Heuvel et al., 2017). Thus, instead of using a proportional threshold, we reduced the complete graphs, deriving thresholds from "shuffled" adult-child pairs. By adapting subsampling-based reduction (Drakesmith et al., 2015) and deconfounding (Epstein et al., 2012), shuffled adult-child pairs were constructed by permuting adults and children who participated in the same experimental condition but independently of each other. This enabled us to control for confounding variables during graph reduction by defining the exchangeability for blockwise permutation, that is, under which conditions the data could be permuted. Specifically, each shuffled signal pair was formed from two signals from an adult and child who did not participate in the same experiment. We generated all possible shuffled adult-child pairs and calculated their connectivity. With this blockwise permutation, assuming exchangeability of the participant IDs while holding the channel combination and condition fixed, we created a set of null distributions. Subsequently, the 95% quantile of each null distribution was set as the threshold for the respective channel and condition. Through the consideration of the 95% quantile of each null-distribution, this blockwise permutation of shuffled pairs endowed the choice of thresholds with a more solid theoretical foundation and generated multiple large null distributions for each fixed variable combination. Setting the channel combination as fixed allowed us to control for systematic differences between channels, such as differences in the signal-to-noise ratio related to the distance between the brain and skull or the amount of hair. Furthermore, setting the experimental condition as fixed allowed us to control for systematic differences related not to social interaction per se but rather, for instance, to differences in activation patterns (see also Kruppa et al., 2021). In the following the reduced graphs are denoted by $G(V_1 \cup V_2, E^-)$, with $G_{complete}(V_1 \cup V_2, E), s.t. E^-(G) \subseteq E(G_{complete})$.

2.3.4 | Graph topology

A graph can be described by a multitude of metrics such as centrality or degree measures. Here we derived both global and nodal density indices from the reduced graphs. While other topological properties of graphs are possible, density is an essential attribute and topological property that is interpretable at both the global and nodal level across different graph types and allows correction for excluded channels.

Global interbrain density, denoted by d_{global}, was defined as the number of interbrain links that survived permutation, normalized with respect to the total maximum number of possible links (Rubinov & Sporns, 2010) after "noisy channels" were excluded.

Nodal interbrain density, denoted by d_{nodal}, was defined as the number of surviving edges for each node of the adult and child participants, again normalized with respect to the total number of possible edges. Thus, the nodal density estimates how strongly the neural signals of a given node are coherent with the signals of all other nodes of the partner and thereby allows the individual contributions of a particular node to the global connectivity to be determined.

## 2.3.5 | Bipartite Erdős-Rényi graphs

To study whether the nodal characteristics differed between the original graphs and graphs with a random nodal structure, we adopted bipartite versions of binomial Erdős-Rényi graphs (Batagelj & Brandes, 2005). Comparing the nodal density distribution of these random graphs with the nodal density distribution of interbrain networks provides an indication of whether it is conceptually valid to study interbrain networks. Therefore, our adjusted generation procedure creates a rewired undirected graph for each original graph such that the generated binomial bipartite Erdős-Rényi graph follows the same global density as the interbrain network and only the nodal properties are disturbed.

In this rewiring process, we define the density-corrected bipartite Erdős-Rényi graph $G_{Reny}(n, m, p)$ with $n = |V_1|, m = |V_2|$, and $p = d_{global}(G(V_1 \cup V_2, E^-))$. The graph $G_{Reny}(n, m, p)$ is constructed with the binomial bipartite Erdős-Rényi algorithm such that the vertex sets $E^*(G_{Reny}) \subseteq E(G_{complete})$, with $G_{complete}(V_1 \cup V_2, E)$ have a corresponding interbrain network $G(V_1 \cup V_2, E^-)$. The probability p of generating an edge in the set $E^*(G_{Reny})$ is set equal to the global density of the corresponding interbrain network $G(V_1 \cup V_2, E^-)$. Hence, this rewiring process aims to preserve the characteristics of the original graph such that $G_{Reny}(n, m, p)$ will have the same global effects as the corresponding interbrain network $G(V_1 \cup V_2, E^-)$.
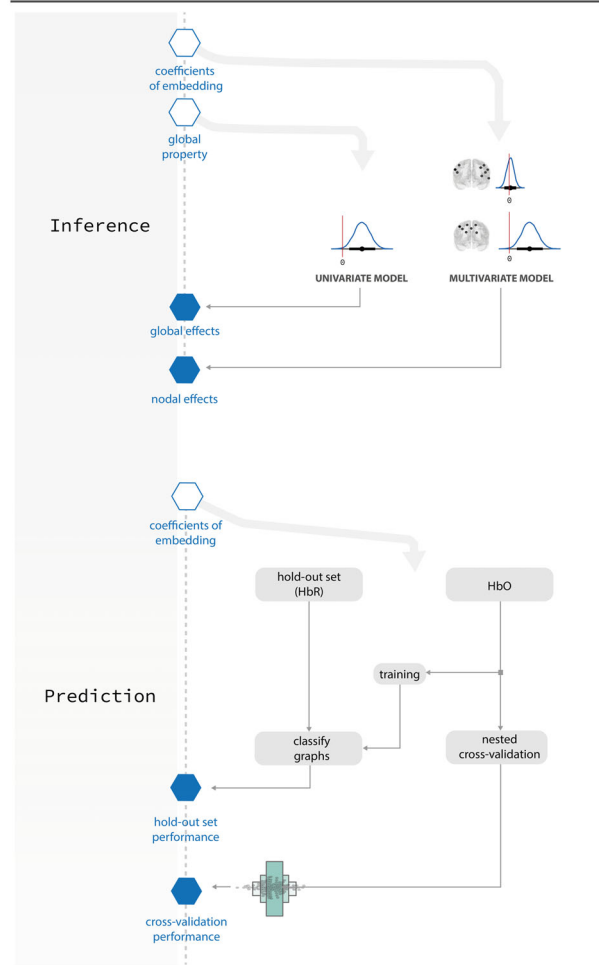
2.4 | Nonnegative matrix factorization

To jointly learn a low dimensional representation across all graphs, NMF was applied to the vectorized nodal metrics. To this end, we

Interbrain Networks
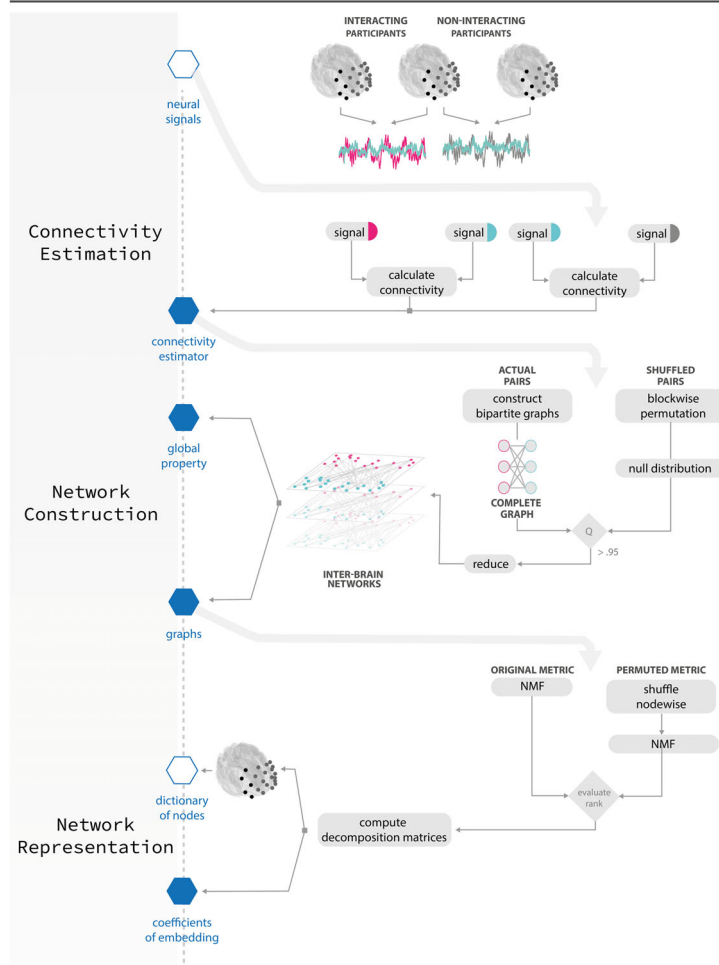
Analysis

FIGURE 2 **Framework for inference and prediction based on interbrain networks.** The proposed framework governs the construction (left) and analysis (right) of interbrain networks. For each dyad and condition, a bipartite graph was constructed from the functional connectivity estimators. The complete bipartite graphs for each dyad were reduced via blockwise permutation using distributions obtained from noninteracting dyads ("shuffled pairs") to reduce the influence of spurious edges and to stratify for confounds. Next, the nodal densities of the interbrain networks were decomposed via NMF to obtain an interpretable, lower-rank representation (embedding). Analysis of the interbrain networks in both the inference and prediction regimes was demonstrated (right). For inference, BHM analysis was introduced, which extends the point estimates of classical frequentist approaches by quantifying the uncertainty of effects using a probability term. First, the global densities of the individual graphs were analyzed to quantify partner and task effects (global effects). Second, the coefficient matrix of our NMF-based network representation was used to formulate a multivariate BHM to yield a more detailed and region-specific quantification of the partner and task effects (nodal effects). Finally, we demonstrated prediction based on interbrain networks using nested-stratified cross-validation of HbO signals and an HbR hold-out set. The performance of our NMF-based network representation was compared with that of state-of-the-art network embeddings and with that of prediction based directly on the functional connectivity estimators

concatenated the nodal densities ($m = 44$ nodes, 22 each for the child and adult) of all graphs ($n = 330$) to form the matrix $V \in R_+^{m \times n}$. As the nodal density ranges between zero and one, V contains only nonnegative entries. This matrix was then approximately factorized into a basis matrix $W \in R_+^{m \times r}$ and a coefficient matrix $H \in R_+^{r \times n}, s.t. V \approx WH$ (Lee & Seung, 1999), where the rank, denoted by r, is constrained to $(n + m)r < nm$. The coefficient matrix H encodes the nodal densities for each of the r components to be used in subsequent analyses. The basis matrix W provides the assignment of nodes to components, thus, can be understood as a "dictionary" in which to look up which nodes

contribute how strongly to a component. Due to the nonnegative constraint, the individual components are strictly additive, making them naturally interpretable, with higher values indicating higher INS. Furthermore, each node can contribute to several components with different weights. These properties set the NMF apart from other methods that yield low vector representations, such as principal component analysis. Of note is that the factorization can also be adjusted to incorporate nodal-in density and nodal-out density metrics of directed interbrain networks, thereby preserving the directionality information of directed connectivity estimators (described in more detail in Text S6).

To obtain stable results, we performed 10,000 iterations of each NMF procedure. The aim of NMF is to minimize the reconstruction error, which generally describes the differences between the original matrix and the product of the decompositions *W* and *H*. The reconstruction error was calculated based on the Frobenius norm. Moreover, the NMF algorithm was initialized via nonnegative double singular value decomposition to ensure a sparse and meaningful representation (Boutsidis & Gallopoulos, 2008). For further details see Text S5.

2.5 | Inference for interbrain networks

To analyze task and partner effects on global density (Section 3.1), a BHM was calculated with the global density as the dependent variable and competition (0 = baseline, 1 = competition), cooperation (0 = baseline, 1 = cooperation), and partner (0 = stranger, 1 = mother) as well as the two-way interactions between task (competition/cooperation) and partner as predictors. Moreover, the empirical findings were validated by formulating a BHM comparing global density of actual and shuffled pairs per condition. To this end, pair (0 = shuffled, 1 = actual), competition (0 = baseline, 1 = competition), cooperation (0 = baseline, 1 = cooperation), partner (0 = stranger, 1 = mother) as well as their two-way interactions were included as predictors. To obtain one shuffled density value for each dyad and condition, global density of all independent ("shuffled") pairs was calculated and the mean values of these distributions per dyad and condition were derived (Kruppa et al., 2021). This ex-post validation verifies whether the graph reduction controls for synchrony of noninteracting pairs.

For the analysis of nodal effects (Section 3.3), an equivalent multivariate BHM was fitted, which included all four NMF components as response variables. The unconditional main effects of competition, cooperation, and partner are reported across partner and task conditions, respectively. In the results section, for each effect of interest, we reported the mean of its estimated marginal posterior distribution as well as its two-sided 90% CI, also termed the "highest posterior density interval," which is defined as a probabilistic interval that is believed to contain a given parameter (Aczel et al., 2020). A detailed description of the model implementations and quality checks can be found in Text S7-S8.

While the BHM benefits from its flexibility, its concrete model formulation can be adapted to a given set of prior information of the specific data. For example, the distribution of alternative topological graph metrics (Section 2.3.4) might make a different linking function of the BHM more appropriate.

2.6 | Classification of interbrain networks

To explore the interbrain network's capacity to predict individual social characteristics, we calculated features based on state-of-the-art network embeddings, vectorized connectivity estimators and our NMF-based embedding. Common machine-learning classifiers, were applied to compare the performance of the embeddings.

2.6.1 | Embeddings and classifiers

State-of-the-art embeddings based on four different methodological approaches were implemented. All embeddings were trained in an unsupervised fashion on graphs constructed as described in Section 2.3.

A. Skipgram-based embeddings
Skipgram refers to a feedforward neural network architecture that was originally applied in natural language processing. Graph2Vec (Narayanan et al., 2017) is a skipgram-based embedding that decomposes each graph into nonlinear substructures by extracting a subgraph for each node via a Weisfeiler-Lehman kernel. These subgraphs are then used in a skipgram model to learn a lower-rank representation in which graphs composed of similar root subgraphs have similar vector representations. Additionally, we applied "Graph and Line graph to vector" (GL2Vec; Chen & Koga, 2019), which uses the edge-adjacency matrix to reformulate the underlying graphs as line graphs and allows the incorporation of additional attributes of nodes as features. We set the bipartite color, that is, which nodes belong to the child and which belong to the partner, as a feature.

B. Spectral-based embeddings
Spectral-based embeddings attempt to generalize convolutional operations to non-Euclidean graph domains. Here, we applied the geometric scattering embedding (Gao et al., 2019), which performs (signal-) processing on graphs (Bronstein et al., 2017) via the wavelet scattering transform. Accordingly, diffusion wavelets calculated from the adjacency matrix and degree matrix are applied to the node signals. These node signals encode a series of node attributes that are captured during random walks across the graph. Again, we adopted the bipartite color as a node attribute.

C. Structural role-based embedding
The local degree profile (LDP; Cai & Wang, 2018) is a simple embedding that is calculated based on five topological features: the nodal degree and the mean, minimum, maximum, and standard deviation of the nodal degree across a 1-neighborhood. These nodal features are mapped into $n = 32$ bins of uniform width. These bins encode each graph into an n-dimensional vector.

D. Our network representation
As described in Section 2.4, we generated a matrix, representing the nodal densities per node and dyad, which was then decomposed via NMF to obtain a low-rank representation of the reduced bipartite graphs. The resulting coefficient matrix was log-transformed, due to its heavy-tailed distribution.

Due to graph reduction (see Section 2.3.3), deconfounded graphs may contain isolated nodes, that is, nodes without any edges. However, most kernel-based and neural-network-based embeddings require a connected graph to ensure that each node of the graph can be reached (weak connectivity assumption). Hence, embeddings A and B were computed based on the connected weighted bipartite

graphs (see Section 2.3.2) but with adjusted connectivity estimators to ensure a fair comparison. The connectivity estimators were block centered with respect to the mean of the shuffled pairs in the corresponding experimental condition, task block and channel combination (see also Section 2.3.3) to mitigate the bias induced by noninteractional synchrony. For further details, see Text S9.

The embeddings were then input into classifiers. Here, we applied a L2-regularized logistic regression classifier (ridge classification: Hoerl & Kennard, 1970), a Gaussian process classifier (Williams & Rasmussen, 2006), and a support vector machine with a linear kernel (Cortes & Vapnik, 1995). We chose these classifiers since they are well established and represent different linear classification approaches.

2.6.2 | Hyperparameter optimization

Hyperparameters indirectly adjust the parametrization of elements of the machine learning algorithms, such as the regularization term of a cost function. Hence, the choice of hyperparameters can strongly influence the outcome of machine learning algorithms (Bergstra & Bengio, 2012). Here, we optimized the hyperparameters separately for each instance of the cross-validation procedure (Section 2.6.3) using Bayesian optimization (Snoek et al., 2012). In the case of a moderate parameter space (e.g., R^d with $d \leq 20$), Gaussian-process-based Bayesian optimization may converge faster to a lower bound solution than alternative approaches. The hyperparameters were optimized based on the area under the receiver operating characteristic curve (AUC-ROC). Further details can be found in Text S9. A detailed parameter space for each classifier and embedding can be found in Table S2.

2.6.3 | Performance evaluation via repeated, stratified, nested k-fold cross-validation

To ensure a generalizable quantification of the predictive classification performance, we applied k-fold cross-validation to the HbO-based graphs as a resampling strategy. In this implementation, the feature matrix F was randomly split into $k = 5$ mutually exclusive subsets of F of approximately equal size, termed folds (for the choice of k see also Kohavi, 2015). The folds denoted by $F_1, F_2, ..., F_k$, were stratified to preserve the class proportions within each fold. For each fold, preprocessing was applied (see Table S2). The hyperparameter optimization was performed separately in a nested loop via stratified cross-validation ($k_{nested} = 3$) within each fold of the outer cross-validation. This nested scheme could avoid leakage effects of the training and test sets caused by the optimization process (Cawley & Talbot, 2010), which we empirically confirmed (Figure S1). For each instantiation, the classifier was trained and optimized on the disjoined set $F \setminus F_t$ of the fold F_t and the feature matrix F and tested on F_t. For each F_t, three performance measures, namely, ROC-AUC, precision and recall (also termed sensitivity), were calculated (see Figure 4 and S2). The

cross-validation procedure was repeated 10 times ($n = 10$) with shuffling to overcome potential bias from the random-partitioning in the cross-validation. The means (μ) and standard deviations (σ) of the performance metrics for each repetition and fold were compared across the test data.

2.6.4 | Out-of-sample performance evaluation via the hold-out set

In contrast to the cross-validation procedure, in which the classifiers were trained based on subsets of graphs constructed from the HbO signals, we also trained classifiers based on all graphs constructed from the HbO data ($n = 330$). Next, these trained classifiers were used for prediction based on the graphs constructed from the HbR signals ($n = 330$). This hold-out set was strongly isolated from all previous training procedures. The postprocessing of the embeddings was analogous to that applied in the cross-validation procedure. The classification performance was measured using the same metrics as in the cross-validation procedure. It should be noted that while this out-of-sample prediction procedure benefits from more training and testing samples compared to the cross-validation procedure, the classifiers also suffer from the challenges arising from prediction across data sets. Importantly, HbO and HbR signals have inverse characteristics, and the magnitude of HbR signals is lower, making this classification task challenging.

3 | RESULTS

3.1 | Inference based on global graph properties

To study interbrain networks and their global properties, we first formulated a complete bipartite graph for each dyad by calculating the pairwise statistical dependencies between neural signals from adult and child regions and reduced the graph via a block permutation-based procedure. The resulting bipartite graphs control for systematic differences not related to the dyad's social interaction. The global density was then quantified as the average number of edges connected to a node, after correction for noisy fNIRS channels, which were excluded from the analyses. The global densities of the resulting graphs were analyzed within a single BHM with competition (0 = baseline, 1 = competition), cooperation (0 = baseline, 1 = cooperation), and partner (0 = stranger, 1 = mother) as well as the two-way interactions between task (competition/cooperation) and partner as predictors.

The results are presented in Figure 1 and Table S3 (Model 1) via the posterior mean (μ), the two-sided 90% credible interval (CI), and the percentage of samples above zero of the marginal posterior distribution. We found evidence for a competitive task effect and, to a weaker degree, evidence for a cooperative task effect, with higher density in the two task conditions compared to the baseline condition (baseline vs. competition: $\mu = 0.13$, CI = [0.02, 0.25], 96.93% > 0; baseline vs. cooperation: $\mu = 0.08$, CI = [−0.02, 0.18], 91.38% > 0).

Further, strong evidence was found for a partner effect, with higher density for mother–child dyads compared to stranger-child dyads ($\mu = 0.13$, CI = [0.04, 0.23], 99.3% > 0). In addition to the stratification for shuffled pairs during the graph reduction procedure, we performed an ex-post validation using shuffled pairs to assess whether neural synchrony of the actual dyads was higher than neural synchrony of independent participants involved in the same experimental condition as described in Section 2.5. The results, described in Table S3, confirmed the previous findings. We found evidence for an interactive effect between pair and partner ($\mu = 0.02$, CI = [0.00, 0.05], 95.58% > 0) and between pair and competition ($\mu = 0.03$, CI = [0.00, 0.06], 95.70% > 0) but only weak evidence for an interactive effect between pair and cooperation ($\mu = 0.02$, CI = [−0.01, 0.04], 88.61% > 0). Directly comparing actual and shuffled pairs, our results showed higher density of actual compared to shuffled pairs for mother–child dyads ($\mu = 0.03$, CI = [0.01, 0.05], 99.08% > 0) but not for stranger–child dyads ($\mu = 0.00$, CI = [−0.01, 0.02], 61.86% > 0) as well as for competition ($\mu = 0.03$, CI = [0.01, 0.05], 98.49% > 0), to a weaker degree for cooperation ($\mu = 0.02$, CI = [−0.00, 0.03], 93.61% > 0) but not for video ($\mu = −0.00$, CI = [−0.02, 0.02], 44.33% > 0).

3.2 | Nodal properties of interbrain networks

In the previous section, we described the global effects of interbrain networks. Building on this, nodal graph metrics may yield further insights into the specific brain regions that facilitate such effects. Since the nodal degree distribution in a complex system has been shown to follow a more heavy-tailed distribution than that in a random graph with uniform edge probabilities (Barabási & Albert, 1999;

Bullmore & Sporns, 2009), we compared the nodal density distributions of the interbrain networks to those of random bipartite Erdős-Rényi graphs. To rule out the possibility that differences these distributions could arise due to differences in global density, we adopted the bipartite Erdős-Renyi graphs so that their rewiring regime preserved the global density of the interbrain networks. In the following, the graphs are compared across all experimental conditions, while separate analyses for the partner and task effects can be found in Text S10, Table S4, Figure S3.

Descriptive results showed that the interbrain networks have both a higher proportion of low-density nodes (density ≤0.01; 83.42% vs. 81.05%) and a higher proportion of high-density nodes (density ≥0.5; 3.34% vs. 1.49%) than Erdős-Rényi graphs with the same global density. These higher proportions of both low-density and high-density nodes indicate that the interbrain networks are more heavy-tailed (Figure 3). These differences between the nodal distributions of the interbrain networks and Erdős-Rényi graphs were confirmed by a significant result from the two-sided Kolmogorov–Smirnov test (0.117, $p < 0.001$). To summarize, our results indicate that the distributions of interbrain graphs were more heavy-tailed than distributions of bipartite Erdös-Renyi graphs with the same global properties.

3.3 | Inference based on nodal graph properties

To analyze the nodal densities while avoiding multiple comparison issues, we derived subsets of nodes showing similar behaviors, reducing the dimensionality of the nodal metrics of each deconfounded graph to four components via NMF. The NMF procedure yields two matrices. The basis matrix can be understood as a dictionary in which

FIGURE 3 Interbrain networks show nonrandom characteristics typical of many real-world networks. The depicted empirical complementary cumulative distribution functions (ECCDFs) show the tail behavior of the nodal density distributions of interbrain networks (blue) and bipartite Erdős-Rényi graphs (black). The ECCDF describes the probability with which a node reaches a specific density. As illustrated, the higher steepness of the interbrain ECCDF curve compared to the bipartite Erdős-Rényi ECCDF curve implies a heavier-tailed nature of the distribution, meaning larger proportions of high- and low-density nodes. Such heavy-tailed characteristics are known to be a fundamental property of many real-world networks
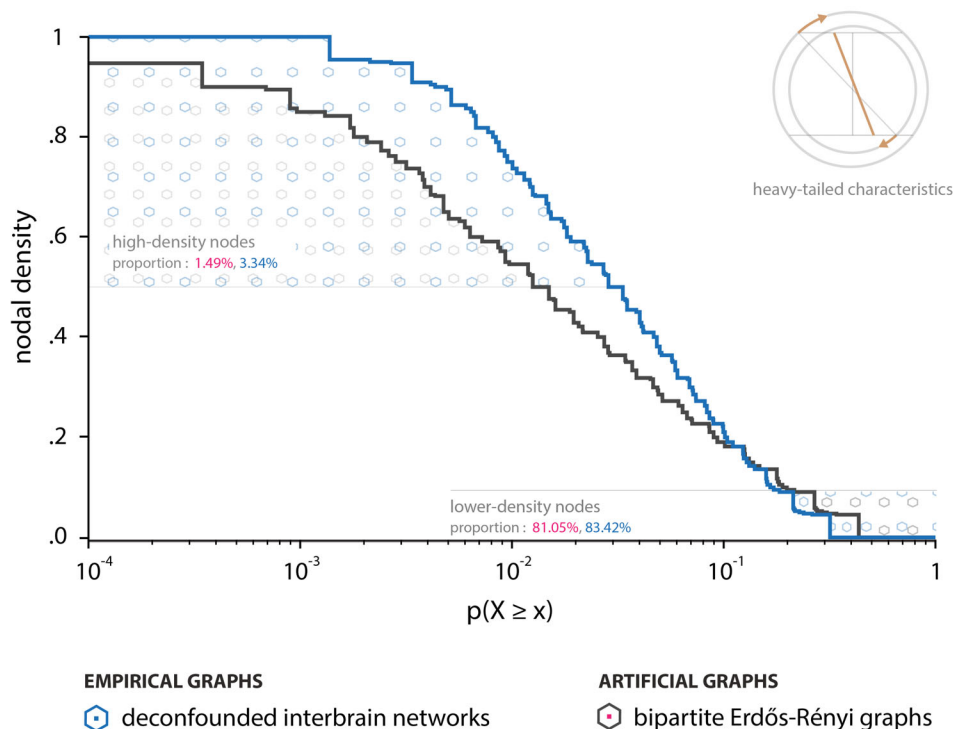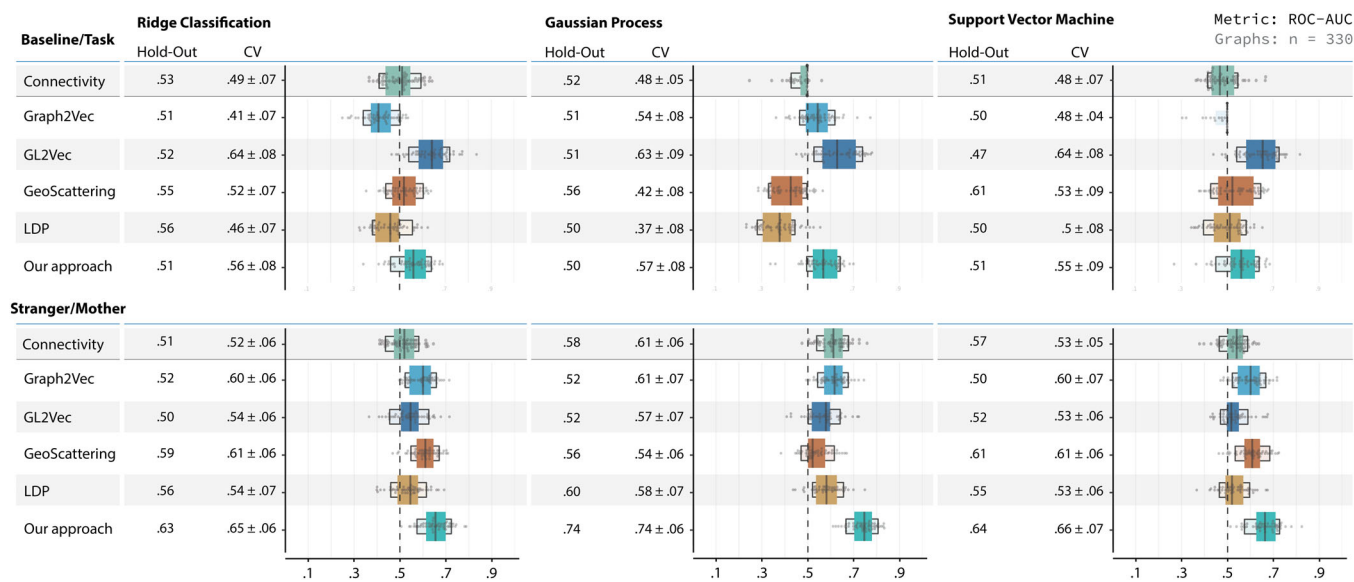
to look up the contributions of each node to each of the components that are constant across dyads and conditions (Figure 1), whereas the coefficient matrix encodes the nodal densities of each dyad in the respective condition. To analyze the nodal densities, a multivariate BHM was again calculated to examine the effects of cooperation, competition, and partner as well as the two-way task × partner interactions on the four NMF components.

The nodal density results confirmed the global results, showing evidence for partner, competition and cooperation effects (Figure 1 and Table S3, Model 2). Specifically, NMF component 1 showed strong evidence for both task effects, with increased INS for cooperation ($\mu = 0.12$, CI = [0.02, 0.22], 97.7% > 0) and competition ($\mu = 0.16$, CI = [0.04, 0.29], 98.56% > 0), as well as weaker evidence for a partner effect ($\mu = 0.08$, CI = [−0.02, 0.18], 91.31% > 0), with increased INS for mother–child dyads compared to stranger-child dyads. Component 2 showed evidence for task × partner interactions (competition × partner: $\mu = -0.27$, CI = [−0.49, −0.05], 2.07% > 0; cooperation × partner: $\mu = -0.21$, CI = [−0.44, 0.02], 6.32% > 0), indicating that mother–child dyads exhibited a higher density than stranger-child dyads in the baseline condition. Component 3 did not show any strong task effects but showed evidence for a partner effect, with higher synchrony of mother–child dyads ($\mu = 0.09$, CI = [0.00, 0.18], 95.58% > 0). Component 4 showed a strong partner effect ($\mu = 0.15$, CI = [0.04, 0.27], 98.44% > 0), a competition and a cooperation effect (baseline vs. competition:

$\mu = 0.13$, CI = [0.02, 0.25], 97.48% > 0; baseline vs. cooperation: $\mu = 0.09$, CI = [−0.02, 0.20], 90.92% > 0). The brain regions that contribute most to each of the components can be found in Figure 1. The obtained NMF solution was validated by showing that (i) the brain regions that contribute to the NMF components show the same or similar effects as the NMF components themselves (Text S11) and (ii) the procedure is not permutation biased, an essential requirement in the graph space, meaning that the results are independent of the ordering of the nodes (Text S12).

3.4 | Exploring predictive applications

As described above, the NMF components were first used in an inference task. Next, we explored whether our network representation can also be applied for a predictive goal. To this end, we trained and tested classifiers to discriminate between graphs according to the task conditions (baseline/task), whereby cooperation and competition are summarized to one category ("task") and compared to baseline, and the partner conditions (stranger/mother). To establish a rigorous evaluation setting given the sample size and the available chromophores, the models were benchmarked via out-of-sample classification performance in the following two regimes (see also Poldrack et al., 2020):

| Baseline/Task | Ridge Classification | | | Gaussian Process | | | Support Vector Machine | | | Metric: ROC–AUC |
	Hold-Out	CV		Hold-Out	CV		Hold-Out	CV		Graphs: n = 330
Connectivity	.53	.49 ± .07		.52	.48 ± .05		.51	.48 ± .07		
Graph2Vec	.51	.41 ± .07		.51	.54 ± .08		.50	.48 ± .04		
GL2Vec	.52	.64 ± .08		.51	.63 ± .09		.47	.64 ± .08		
GeoScattering	.55	.52 ± .07		.56	.42 ± .08		.61	.53 ± .09		
LDP	.56	.46 ± .07		.50	.37 ± .08		.50	.5 ± .08		
Our approach	.51	.56 ± .08		.50	.57 ± .08		.51	.55 ± .09		
Stranger/Mother										
Connectivity	.51	.52 ± .06		.58	.61 ± .06		.57	.53 ± .05		
Graph2Vec	.52	.60 ± .06		.52	.61 ± .07		.50	.60 ± .07		
GL2Vec	.50	.54 ± .06		.52	.57 ± .07		.52	.53 ± .06		
GeoScattering	.59	.61 ± .06		.56	.54 ± .06		.61	.61 ± .06		
LDP	.56	.54 ± .07		.60	.58 ± .07		.55	.53 ± .06		
Our approach	.63	.65 ± .06		.74	.74 ± .06		.64	.66 ± .07		

FIGURE 4 Network representations of interbrain networks might improve the classification of interaction partners and tasks. Classification performance based on network representations and functional connectivity was compared using three widely applied classifiers. For the two classification tasks, baseline/task (top) and stranger/mother (bottom), out-of-sample performance was evaluated via the ROC-AUC. The evaluation regimes consisted of nested cross-validation ($CV; k_{out} = 5, k_{inner} = 3$) and an additional hold-out set (HbR). The CV was stratified to preserve the class distribution across folds and was repeated ($n = 10$) to avoid the bias that may arise from the splitting of the folds. The "letter-value plots" and the table display the means, standard deviations, and quantiles of the CV performance. To quantify each embedding's ability to generalize across data sets, models were trained on the HbO data and used to predict task and partner conditions based on the HbR data. The obtained ROC-AUC results for the CV and hold-out set regimes showed that our proposed NMF-based network representation can successfully discriminate between partners (stranger/mother). GL2Vec showed better performance for task prediction (baseline/task), although task prediction generally appeared more challenging, particularly on the hold-out data. In addition to the network embeddings, the same analysis was performed directly on the functional connectivity estimators. The results showed that prediction based on network representations generally outperformed prediction based on functional connectivity, supporting the validity of using interbrain networks for the analysis of INS

i. A repeated, stratified, nested k-fold cross-validation procedure was applied based on graphs derived from the HbO signals (see Section 2.6.3).

ii. In a true hold-out regime, the classifiers were trained on all HbO graphs and tested on graphs derived from the HbR signals (see Section 2.6.4).

First, we assessed whether the predictive models were able to discriminate task conditions (baseline/task) in the two evaluation settings (i, ii). For (i), GL2Vec and our network representation showed results that were better than chance across all classifiers (Figure 4). However, neither method was able to uphold its performance on the hold-out set. GeoScattering was superior in regime (ii) but showed performance below chance on (i) for the Gaussian process classifier. Prediction based on the functional connectivity estimators did not reach the level of chance for any classifier in both regimes (i) and (ii).

Second, we evaluated whether the graphs could be classified based on the partner condition (stranger/mother). Consistent with the partner effect evidenced by the BHM analysis (see Sections 3.1 and 3.3), most embeddings showed results better than chance. Generally, the performance was higher for partner classification than for task classification across all classifiers, with the exception of GL2Vec. Our NMF-based approach outperformed all network embeddings as well as the functional connectivity estimators across all classifiers in both evaluation settings (i, ii). Again, GeoScattering reached competitive results in regime (ii). Both GeoScattering and our approach showed comparable performance between (i) and (ii). GL2Vec showed relatively consistent performance better than chance in the cross-validation regime (i) but exhibited poor generalization across data sets (ii) (Figure S2B). A closer look at the recall of Graph2Vec, revealed that it tended to assign a majority of graphs to the same class (Figure S2B). Detailed results, including additional performance metrics, classifiers, and parameters, can be found in the Figure S2 and Table S2.

4 | DISCUSSION

Extending the isolated view of a single brain by examining the statistical dependencies between interacting brains, may shed new light on the neurobiological underpinnings of social interactions. While network analysis is an established, state-of-the-art tool for investigating brain connectivity, hyperscanning studies often analyze only the pairwise relationships between the brain regions of one participant and the brain regions of another. Moving beyond this pairwise analysis, we here present the first comprehensive analytical framework for inference and prediction based on bipartite interbrain networks.

4.1 | Bipartite graphs as a representation of interbrain synchrony

In hyperscanning studies, INS is often calculated region-by-region between homologous brain regions (e.g., Cui et al., 2012; Djalovski et al., 2020; Reindl et al., 2018). However, this approach implicitly

assumes that synchrony arises primarily in homologous brain regions, thereby neglecting the fact that each brain is organized and operates as a complex network. Along these lines, mentalizing in more natural social interactions may elicit activity in a broader network than the activity elicited in response to static, isolated stimuli. This may be the case either because this broader network may be characteristic of how the brain processes real-world social stimuli or because multiple processes are simultaneously engaged (Redcay & Moraczewski, 2020). In the cooperative and competitive tasks, these processes may include prediction and adaption processes, emotional process, social comparison processes but also processing of stimuli and shared attention to the task (see also Kruppa et al., 2021; Reindl et al., 2018). Thus, information processing in social interactions is complex and likely goes beyond activity occurring in a single, circumscribed brain region, but most likely engages multiple, interacting brain networks in both subjects. Examining synchrony only in a pairwise, region-by-region fashion may fall short of capturing this complexity. This general critique of pairwise analyses is grounded in system theory, as such analyses ignore the multiple dependencies among the entities of a system (Goyal & Ferrara, 2018).

In addition to neuroscientific arguments for a graph-based approach, there are also several methodological advantages. First, research shows that interindividual differences exist in both the brain structures and networks in which cognitive, emotional, and social processes are substantiated (e.g., Kanai & Rees, 2011; Vanderwal et al., 2017). These differences may be even more pronounced when studying dyads of different age groups, such as adults and infants/children, considering the ongoing maturation of brain structures and development of functions in the latter group (see also Power et al., 2010). Such inter-individual differences may obscure effects in a pairwise (region-by-region) analysis. On top of such anatomical differences, uncertainties in channel localization may arise due to the spatial resolution of the technique, as both fNIRS and EEG are inferior in this respect to functional magnetic resonance imaging (fMRI).

Yet, it should be noted that the most appropriate choice of analytical method depends on the specific research question at hand. For example, for a research question testing hypotheses on the pairwise relationship between specific brain regions, it may be beneficial to calculate a connectivity estimator between these (e.g., homologous) brains regions. Since a more complete description of a complex system may come at the expense of precision on entity level, accounting for multiple dependencies across regions that are not of research interest may not be advantageous. For other research questions targeting multiple dependencies, inferential models that aim to mimic the underlying biophysical processes might be beneficial, such as the dynamic causal model (Bilek et al., 2022). Alternatively, other research questions could favor avoiding such priors via the proposed graph-based framework when focusing on global and local insights within the cross-brain network.

To summarize, we propose to move away from a pairwise analysis and towards a graph-based analytical approach when measuring INS between interacting persons for the reasons described above, namely, (i) the engagement of brain networks rather than single brain regions, (ii) inter-individual anatomical differences, and (iii) uncertainties in

localization. Yet, this system perspective comes with several analytical challenges, such as how to identify and deal with spurious edges (see Sections 2.3.3 and 3.1) and how to gain local insights into a network (see Section 3.3).

4.2 | Nodal properties of interbrain networks

Investigating the properties of interbrain networks, we have shown that the nodal density distribution is more heavy-tailed than the distribution of random graphs with the same global density. Such skewed distributions can be found on nearly all physiological and anatomical levels of the brain as well as in many networks—encountered in fields ranging from biology to social sciences to infrastructure (Broido & Clauset, 2019; Buzsáki & Mizuseki, 2014). Thus, these results indicate that interbrain networks have properties similar to those of many complex real-world networks and are significantly different from random graphs. Such high-density nodes may act as hubs and play an important structural role in their respective network topologies (see also Henderson et al., 2012; Rossi & Ahmed, 2014). These findings have important practical implications for future research. First, many classical frequentist tests require a Gaussian distribution and thus are ill suited for analyzing heavy-tailed distributions. Second, this study provides the first evidence that some nodes have unproportionally high degrees, which implies that focusing on global metrics alone is not sufficient to examine how brain activities synchronize during social interactions. In contrast, a nodal analysis not only may yield more topological details but could also reveal more localized effects that would otherwise be obscured when considering global metrics alone. Finally, these findings should be considered a first starting point that raises new questions about the role of high-density regions in interacting yet structurally unconnected brains. For instance, future research may address whether and how high-density nodes in inter- and intrabrain networks are related to each other.

4.3 | Inference based on interbrain networks

We have successfully demonstrated that interbrain networks can be constructed in the form of bipartite graphs. Regarding the global density, our BHM results show that the density was higher for the two interactive conditions in our study (cooperation/competition) than for the noninteractive baseline condition. In line with our previous findings in male children and adolescents (Kruppa et al., 2021), this indicates that interbrain networks can encode interaction specific neural processes distinct from those associated with a noninteractive, rest-like state. Moreover, the mother–child dyads showed higher INS than the stranger-child dyads. This was a general effect, observed across the baseline and task conditions, suggesting that the brain responses may be more in tune with each other due to genetic factors, the dyad's affiliation or shared past experiences. This finding is consistent with several other studies showing increased INS in close relationships when performing joint motor tasks (Djalovski et al., 2020;

Kruppa et al., 2021; Pan et al., 2017; Reindl et al., 2018). For a more detailed discussion of the task and partner effects, see (Reindl et al., 2022).

Subsequently, a NMF was used to enable the comparison of nodal properties. The NMF is applied in various domains and to address a variety of problems, such as the detection of communities or the clustering of neural connections (e.g., Goldstein et al., 2018). One of its key advantages is its ability to extract sparse and meaningful features. While there are nearly infinite possibilities to decompose a matrix or picture, NMF allows to identify latent factors that can be interpreted as "parts-of-whole", such as eyes within a picture of a person (Lee & Seung, 1999).

Analogously, here, we found that the NMF was able to decompose the global density effect into subsets of nodes showing particular distinct result patterns. Specifically, we found cooperative and competitive task effects in components 1 and 4, located mostly in the left and right lateralized dorsolateral and frontopolar brain regions of adult and child, as well as evidence for a partner effect in components 3 and 4, located mostly in the left, right and medial frontopolar brain regions of both partners as well as the right dorsolateral brain regions of the adult (for further discussion of the brain regions, see Reindl et al., 2022). Thus, the NMF results confirmed the global results but provided additional topological details and thereby allow to interpret the brain regions which drive neural synchrony. Importantly, we were able to show that our NMF based method is not permutation biased and thus fulfills an essential requirement for graph representations.

Since we applied an undirected connectivity estimator, the results assume a bidirectional interaction between signals (Section 2.3.1). For research inquiries addressing the directionality of INS, for example, from adult to child or child to adult, we have described in Text S6 how the NMF-based embedding can be adapted to encode-directed graphs. Further, the current approach is based on a static graph. A future extension to dynamic graphs could potentially gain additional insights into the temporal unfolding of neural synchrony.

4.4 | Classification performance of network embeddings

Next, we extended the application of the NMF-based method from inference to prediction to answer the question of whether network representations have the capacity to predict experimental conditions on an individual level. Although all network representations achieved a performance above the chance level for predicting the interaction partner (stranger/mother), the NMF-based representation presented here showed the best performance across both a rigorous cross-validation regime and prediction on a second data set (HbR hold-out set). In contrast to our NMF-based representation, graph embeddings often rely on higher-order proximities, that is, similarities that can be derived from complex operations on the adjacency matrix, at the cost that they are difficult or impossible to interpret. This makes them suitable for predictive tasks but less so for inference tasks (Zhang

et al., 2020). In contrast, a few embeddings, such as LDP (Cai & Wang, 2018), exploit the possibility that nodes may play similar structural roles within the network. LDP is therefore most similar to our approach but considers the structural role of a node by its degree, whereas our NMF-based representation is based on nodal density, which is less biased by excluded channels, and considers more complex, additive patterns. LDP showed a ROC-AUC above chance level and yielded somewhat competitive results compared to other embeddings but was consistently outperformed by our NMF-based approach. This finding supports the concept of encoding the brain regions' structural roles but shows that considering more complex, additive patterns and a graph metric that is not biased by excluded channels may be advantageous. Importantly, it was found that both GeoScattering, a state-of-the-art graph embedding based on adjusted wavelet transformation for graphs, and the NMF-based representation could be trained on one data set and successfully ported to another data set, indicating reasonable generalization across diverse settings. Such inhomogeneous data sets are likely to arise in neuroscience due to different experimental designs, possible confounders or reduced standardization in naturalistic settings. Consequently, to elucidate the neural basis of social interaction, generalization across data sets will probably be a challenging key factor for predictive applications.

While partner classification showed promising results, task classification was challenging. GL2Vec showed the best performance and reasonable behavior in terms of precision and recall in the cross-validation setting (Figures 4 and S2A). Interestingly, GL2Vec, which is based on a line graph (Evans & Lambiotte, 2009), achieved relatively consistent performance above the chance level for cross-validation but showed poor generalization across data sets. A line graph provides an edge-centric representation because it maps the edges of the original graph onto the nodes of the line graph. Recent studies (e.g., Faskowitz et al., 2020) suggest that edge-centric approaches may be advantageous in considering individual differences in brain networks. Overall, the lower performance for task classification may be explained by the class imbalance between the baseline ($n_{baseline} = 64$) and task ($n_{task} = 266$) conditions. Very high and low precision and recall values, respectively, indicate that especially Graph2Vec but also our NMF-based approach might suffer from class imbalance and, for some instances, classify the majority of graphs only as task or only as baseline (Figure S2A). Although we have made some effort to tackle the imbalance (Text S9), low performance may be driven by the fact that not all dyads showed an increased INS during task conditions.

A central finding from our analyses pertains to the general capacities of network representations. Network representations performed at least as well as or better than predictions based directly on the functional connectivity estimators. These results indicate that network representations of INS may carry further information compared to the "pure" functional connectivity estimators. Importantly, in contrast to the other embeddings, our NMF-based approach not only yielded competitive results but also preserved interpretability.

4.5 | Limitations

Although the current study provides the first evidence that the proposed graph analytical framework can be successfully applied to studies on interbrain connectivity, a number of caveats must be considered. First, intrabrain relationships were not explicitly modeled in the current study, a limitation that could be addressed by other graph formulations. Second, the current reduction procedure enables stratification for confounds but does not ensure a weakly connected graph without isolated nodes, that is, nodes with no connection to any node. Thus, it cannot be applied to network embeddings that require all nodes to be connected. Third, while our results show that network representations could potentially serve for predictive tasks in hyperscanning, the baseline/task classification task in particular presents a challenge that will require further improvements, which are beyond the scope of this article.

5 | CONCLUSION

Network models may contribute to a better understanding of the neural underpinnings of social interaction, possibly providing a more precise analytical operationalization of the concept of INS than pairwise statistical dependencies. Based on bipartite graphs, we have proposed a novel computational approach to capture the dynamic interplay of brain regions that considers the multiple dependencies across the brain regions of interacting persons. The heavy-tailed characteristics of these interbrain networks, differing from those of random graphs, as well as the replication of previous findings regarding partner and task effects and the demonstrated predictive power of interbrain network representations argue for their validity. This framework comes with several advantages: (i) it yields both global and nodal insights; (ii) through Bayesian modeling, it avoids the problems associated with multiple comparisons, accounts for skewed distributions, and provides uncertainty estimates; and (iii) it is applicable for both inference and prediction tasks while providing interpretable insights. Thus, the proposed framework has potentially far-reaching implications, as it enables the exploration of the specific brain regions in a network that seed neural synchrony between interacting persons—an approach that could also be applied for the topological analysis of intrabrain networks. Above and beyond inferential inquiries, our results are the first to demonstrate that network embeddings could provide an informative representation of interpersonal synchrony and may thereby stimulate further research into prediction based on graph-structured data.

AUTHOR CONTRIBUTIONS

CONFLICT OF INTEREST

The authors declare that they have no competing interests.

DATA AVAILABILITY STATEMENT

Source data are provided in the supplementary materials. The data that support the findings of this study are available from the corresponding author upon reasonable request after confirmation from the ethical committee of our institution. Details of the computational implementation are described in the supplementary materials (see also text S13). The source code of interbrain networks and embedding are available at the corresponding GitHub repository (https://github.com/ChristianGerloff/interbrain-network-embedding).

ORCID

Christian Gerloff https://orcid.org/0000-0003-2718-2089

REFERENCES

Achard, S., Salvador, R., Whitcher, B., Suckling, J., & Bullmore, E. (2006). A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *The Journal of Neuroscience*, *26*, 63–72.

Aczel, B., Hoekstra, R., Gelman, A., Wagenmakers, E.-J., Klugkist, I. G., Rouder, J. N., Vandekerckhove, J., Lee, M. D., Morey, R. D., & Vanpaemel, W. (2020). Discussion points for Bayesian inference. *Nature Human Behaviour*, *1-3*, 561–563.

Babiloni, F., & Astolfi, L. (2014). Social neuroscience and hyperscanning techniques: Past, present and future. *Neuroscience and Biobehavioral Reviews*, *44*, 76–93.

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, *286*, 509–512.

Bastos, A. M., & Schoffelen, J. M. (2016). A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Frontiers in Systems Neuroscience*, *9*, 175.

Batagelj, V., & Brandes, U. (2005). Efficient generation of large random networks. *Physical Review E*, *71*, 036113.

Bergstra, J., & Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, *13*, 281–305.

Bilek, E., Zeidman, P., Kirsch, P., Tost, H., Meyer-Lindenberg, A., & Friston, K. (2022). Directed coupling in multi-brain networks underlies generalized synchrony during social exchange. *NeuroImage*, *252*, 119038.

Boutsidis, C., & Gallopoulos, E. (2008). SVD based initialization: A head start for nonnegative matrix factorization. *Pattern Recognition*, *41*, 1350–1362.

Broido, A. D., & Clauset, A. (2019). Scale-free networks are rare. *Nature Communications*, *10*, 1–10.

Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., & Vandergheynst, P. (2017). Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, *34*, 18–42.

Bullmore, E., & Sporns, O. (2009). Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nature Reviews. Neuroscience*, *10*, 186–198.

Buzsáki, G., & Mizuseki, K. (2014). The log-dynamic brain: How skewed distributions affect network operations. *Nature Reviews. Neuroscience*, *15*, 264–278.

Cai, C., & Wang, Y. (2018). A simple yet effective baseline for non-attributed graph classification. *arXiv*, arXiv:1811.03508.

Cawley, G. C., & Talbot, N. L. (2010). On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, *11*, 2079–2107.

Chami, I., Abu-El-Haija, S., Perozzi, B., Ré, C., & Murphy, K. (2022). Machine learning on graphs: A model and comprehensive taxonomy. *Journal of Machine Learning Research*, *23*, 1–64.

Chen, H., & Koga, H. (2019). GL2vec: Graph Embedding Enriched by Line Graphs with Edge Features. International conference on neural information processing, 3.14.

Ciaramidaro, A., Toppi, J., Casper, C., Freitag, C., Siniatchkin, M., & Astolfi, L. (2018). Multiple-brain connectivity during third party punishment: An EEG hyperscanning study. *Scientific Reports*, *8*, 1–13.

Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, *20*, 273–297.

Cui, X., Bryant, D. M., & Reiss, A. L. (2012). NIRS-based hyperscanning reveals increased interpersonal coherence in superior frontal cortex during cooperation. *NeuroImage*, *59*, 2430–2437.

Djalovski, A., Dumas, G., Kinreich, S., & Feldman, R. (2020). Human attachments shape interbrain synchrony toward efficient performance of social goals. *NeuroImage*, *226*, 117600.

Drakesmith, M., Caeyenberghs, K., Dutt, A., Lewis, G., David, A., & Jones, D. K. (2015). Overcoming the effects of false positives and threshold bias in graph theoretical analyses of neuroimaging data. *NeuroImage*, *118*, 313–333.

Epstein, M. P., Duncan, R., Jiang, Y., Conneely, K. N., Allen, A. S., & Satten, G. A. (2012). A permutation procedure to correct for confounders in case-control studies, including tests of rare variation. *American Journal of Human Genetics*, *91*, 215–223.

Erdös, P., & Rényi, A. (1959). On random graphs I. *Universitatis Debreceniensis*, *6*, 290–297.

Evans, T. S., & Lambiotte, R. (2009). Line graphs, link partitions, and overlapping communities. *Physical Review E*, *80*, 016105.

Falk, E. B., & Bassett, D. S. (2017). Brain and social networks: Fundamental building blocks of human experience. *Trends in Cognitive Sciences*, *21*, 674–690.

Faskowitz, J., Esfahlani, F. Z., Jo, Y., Sporns, O., & Betzel, R. F. (2020). Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture. *Nature Neuroscience*, *23*, 1644–1654.

Feldman, R. (2012). Parent-infant synchrony: A biobehavioral model of mutual influences in the formation of affiliative bonds. *NeuroImage*, *77*, 42–51.

Finn, E. S., Shen, X., Scheinost, D., Rosenberg, M. D., Huang, J., Chun, M. M., Papademetris, X., & Constable, R. T. (2015). Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, *18*, 1664–1671.

Fornito, A., Zalesky, A., & Bullmore, E. (2016). *Fundamentals of brain network analysis*. Academic Press, chap. 4.

Fortin S, (1996) "The graph isomorphism problem" (Tech. Rep. TR96-20, University of Alberta).

Gao, F., Wolf, G., & Hirn, M. (2019). Geometric scattering for graph data analysis. *Proceedings of the 36th International Conference on Machine Learning*, 97, 2122–2131.

Gerloff, C. (2022) Github repository containing "channel2mni" with most likely MNI coordinates of fNIRS channels. https://github.com/ChristianGerloff/interbrain-network-embedding

Goldstein, P., Weissman-Fogel, I., Dumas, G., & Shamay-Tsoory, S. G. (2018). Brain-to-brain coupling during handholding is associated with pain reduction. *Proceedings of the national academy of sciences*, 115, E2528–E2537.

Goyal, P., & Ferrara, E. (2018). Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems*, 151, 78–94.

Hamilton, A. F. D. C. (2021). Hyperscanning: Beyond the hype. *Neuron*, 109, 404–407.

Hari, R., Henriksson, L., Malinen, S., & Parkkonen, L. (2015). Centrality of social interaction in human brain function. *Neuron*, 88, 181–193.

Henderson, K., Gallagher, B., Eliassi-Rad, T., Tong, H., Basu, S., Akoglu, L., Koutra, D., Faloutsos, C., & Li, L. (2012). RolX: structural role extraction & mining in large graphs. *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1231–1239.

Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12, 55–67.

Kanai, R., & Rees, G. (2011). The structural basis of inter-individual differences in human behaviour and cognition. *Nature Reviews. Neuroscience*, 12, 231–242.

Kohavi, R. (2015). A study of cross validation and bootstrap for accuracy estimation and model selection. *Proceedings of the 14th international joint conference on artificial intelligence*, 14, 1137–1143.

Kruppa, J. A., Reindl, V., Gerloff, C., Weiss, E. O., Prinz, J., Herpertz-Dahlmann, B., Konrad, K., & Schulte-Rüther, M. (2021). Brain and motor synchrony in children and adolescents with ASD–an fNIRS hyperscanning study. *Social Cognitive and Affective Neuroscience*, 16, 103–116.

Kruschke, J. (2014). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan*. Academic Press.

Lancaster, J. L., Woldorff, M. G., Parsons, L. M., Liotti, M., Freitas, C. S., Rainey, L., Kochunov, P. V., Nickerson, D., Mikiten, S. A., & Fox, P. T. (2000). Automated Talairach atlas labels for functional brain mapping. *Human Brain Mapping*, 10, 120–131.

Latora, V., & Marchiori, M. (2001). Efficient behavior of small-world networks. *Physical Review Letters*, 87, 198701.

Lee, D., & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401, 788–791.

Narayanan, A., Chandramohan, M., Venkatesan, R., Chen, L., Liu, Y., & Jaiswal, S. (2017). graph2vec: Learning distributed representations of graphs. *arXiv*, arXiv:1707.05005.

Nelson, E. E., Leibenluft, E., McClure, E. B., & Pine, D. S. (2005). The social re-orientation of adolescence: A neuroscience perspective on the process and its relation to psychopathology. *Psychological Medicine*, 35, 163–174.

Pan, Y., Cheng, X., Zhang, Z., Li, X., & Hu, Y. (2017). Cooperation in lovers: An f NIRS-based hyperscanning study. *Human Brain Mapping*, 38, 831–841.

Passingham, R. E., Stephan, K. E., & Kötter, R. (2002). The anatomical basis of functional localization in the cortex. *Nature Reviews. Neuroscience*, 3, 606–616.

Poldrack, R. A., Huckins, G., & Varoquaux, G. (2020). Establishment of best practices for evidence for prediction: A review. *JAMA Psychiatry*, 77, 534–540.

Power, J. D., Fair, D. A., Schlaggar, B. L., & Petersen, S. E. (2010). The development of human functional brain networks. *Neuron*, 67, 735–748.

Redcay, E., & Moraczewski, D. (2020). Social cognition in context: A naturalistic imaging approach. *NeuroImage*, 216, 116392.

Reindl, V., Gerloff, C., Scharke, W., & Konrad, K. (2018). Brain-to-brain synchrony in parent-child dyads and the relationship with emotion regulation revealed by fNIRS-based hyperscanning. *NeuroImage*, 178, 493–502.

Reindl, V., Konrad, K., Gerloff, C., Kruppa, J. A., Bell, L., & Scharke, W. (2019). Conducting hyperscanning experiments with functional near-infrared spectroscopy. *Journal of Visualized Experiments*, e58807.

Reindl, V., Wass, S., Leong, V., Scharke, W., Wistuba, S., Wirth, C. L., Konrad, K., & Gerloff, C. (2022). Multimodal hyperscanning reveals that synchrony of body and mind are distinct in mother-child dyads. *NeuroImage*, 251, 118982.

Rosenthal, G., Váša, F., Griffa, A., Hagmann, P., Amico, E., Goñi, J., Avidan, G., & Sporns, O. (2018). Mapping higher-order relations between brain structure and function with embedded vector representations of connectomes. *Nature Communications*, 9, 1–12.

Rossi, R. A., & Ahmed, N. K. (2014). Role discovery in networks. *IEEE Transactions on Knowledge and Data Engineering*, 27, 1112–1131.

Rubinov, M., & Sporns, O. (2010). Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage*, 52, 1059–1069.

Santamaria, L., Noreika, V., Georgieva, S., Clackson, K., Wass, S., & Leong, V. (2020). Emotional valence modulates the topology of the parent-infant inter-brain network. *NeuroImage*, 207, 116341.

Scholkmann, F., Spichtig, S., Muehlemann, T., & Wolf, M. (2010). How to detect and reduce movement artifacts in near-infrared imaging using moving standard deviation and spline interpolation. *Physiological Measurement*, 31, 649–662.

Scholkmann, F., & Wolf, M. (2013). General equation for the differential pathlength factor of the frontal human head depending on wavelength and age. *Journal of Biomedical Optics*, 18, 105004.

Shmueli, G. (2010). To explain or to predict? *Statistical Science*, 25, 289–310.

Singh, A. K., Okamoto, M., Dan, H., Jurcak, V., & Dan, I. (2005). Spatial registration of multichannel multi-subject fNIRS data to MNI space without MRI. *NeuroImage*, 27, 842–851.

Snoek, J., Larochelle, H., & Adams, R. (2012). Practical Bayesian optimization of machine learning algorithms. *Advances in Neural Information Processing Systems*, 25, 2951–2959.

Toppi, J., Ciaramidaro, A., Vogel, P., Mattia, D., Babiloni, F., Siniatchkin, M., & Astolfi, L. (2015). Graph theory in brain-to-brain connectivity: a simulation study and an application to an EEG hyperscanning experiment. *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (pp. 2211–2214).

Torrence, C., & Compo, G. P. (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society*, 79, 61–78.

Tsuzuki, D., Jurcak, V., Singh, A. K., Okamoto, M., Watanabe, E., & Dan, I. (2007). Virtual spatial registration of stand-alone fNIRS data to MNI space. *NeuroImage*, 34, 1506–1518.

van den Heuvel, M. P., de Lange, S. C., Zalesky, A., Seguin, C., Yeo, B. T., & Schmidt, R. (2017). Proportional thresholding in resting-state fMRI functional connectivity networks and consequences for patient-control connectome studies: Issues and recommendations. *NeuroImage*, 152, 437–449.

van den Heuvel, M. P., & Sporns, O. (2013). Network hubs in the human brain. *Trends in Cognitive Sciences*, 17, 683–696.

Vanderwal, T., Eilbott, J., Finn, E. S., Craddock, R. C., Turnbull, A., & Castellanos, F. X. (2017). Individual differences in functional connectivity during naturalistic viewing conditions. *NeuroImage*, *157*, 521–530.

Williams, C. K., & Rasmussen, C. E. (2006). *Gaussian processes for machine learning* (Vol. 2). MIT Press.

Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Philip, S. Y. (2021). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, *32*, 4–24.

Zhang, D., Yin, J., Zhu, X., & Zhang, C. (2020). Network representation learning: A survey. *IEEE Transactions on Big Data*, *6*, 3–28.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.