

<https://doi.org/10.1038/s42005-025-02472-9>

# The microscale organization of directed hypergraphs

Check for updates

Quintino Francesco Lotito<sup>1,2</sup>✉, Alberto Vendramini<sup>2</sup>, Alberto Montresor<sup>2</sup> & Federico Battiston<sup>1</sup>✉

Many real-world systems exhibit higher-order interactions beyond pairwise links. Such interactions are modeled by undirected hypergraphs where edges can connect any number of vertices, but without capturing the directional nature of many real-world interactions. Directed hypergraphs overcome this limitation by distinguishing source and target sets within each hyperedge, enabling analysis of directional information flow. Here, we provide a framework to characterize the structural organization of directed higher-order networks at their microscale. We extract the fingerprint of a directed hypergraph, capturing the frequency of hyperedges with a certain source and target sizes, and use this information to compute differences in higher-order connectivity patterns among real-world systems. Then, we investigate the overlap among sources and targets to reveal recurring sets of co-sending and co-receiving nodes. We define reciprocity in hypergraphs using exact, strong, and weak definitions to quantify the extent to which hyperedges are reciprocated. Finally, we extend motif analysis to identify recurring interaction patterns and extract the building blocks of directed hypergraphs. We validate our framework on empirical datasets, including Bitcoin transactions, metabolic networks, and citation data, revealing structural principles behind the organization of real-world systems.

Accurately modeling interactions among entities is crucial to understand the properties of many complex systems. Traditional network models focus on pairwise connections between nodes<sup>1,2</sup>, neglecting the complexities of systems where multiple units interact simultaneously. Such higher-order interactions are prevalent in various domains, including social networks<sup>3–5</sup>, folksonomies<sup>6</sup>, ecological systems<sup>7</sup>, chemical reactions<sup>8</sup> including metabolic pathways<sup>9</sup>, and the brain<sup>10,11</sup>.

Hypergraphs<sup>12</sup> provide a framework for explicitly encoding higher-order interactions, representing them as hyperedges connecting multiple nodes simultaneously. By preserving group-based interactions, they improve our ability to understand the structures and dynamics of systems with many-body interactions<sup>13,14</sup>. Recently, a variety of measures have been introduced or extended to capture the higher-order organization of complex systems, including centrality<sup>15,16</sup>, community structure<sup>17–19</sup> and motifs<sup>20–22</sup>. Moreover, new models have allowed to describe systems' evolution<sup>23–25</sup>, and highlight the importance of higher-order interactions in shaping emergent behaviors in diffusion<sup>26,27</sup>, synchronization<sup>28–30</sup>, spreading<sup>31,32</sup> and evolutionary dynamics<sup>33</sup>.

Most research has so far focused on undirected hypergraphs, which fail to capture the directional nature of many real-world interactions. For example, in a metabolic reaction, a set of reactants transforms into a set of products<sup>9</sup>. Similarly, in a Bitcoin transaction, multiple source wallets may

transfer funds simultaneously to multiple target wallets<sup>34</sup>. To accurately encode such interactions, models must incorporate directionality into their representations. In this sense, directed hypergraphs enhance modeling by distinguishing between source and target sets in each hyperedge<sup>35</sup>. Tools to study directed hypergraphs are largely underdeveloped, with notable exceptions in areas, such as null models<sup>36</sup>, synchronization<sup>37</sup>, overlapping patterns between two hyperedges of limited size<sup>38</sup>, and some early proposals to define reciprocity<sup>39,40</sup>.

In this work, we introduce measures and tools to characterize the microscale organization of real-world directed hypergraphs. First, we discuss a decomposition into fundamental interaction types: one-to-one, one-to-many, many-to-one and many-to-many. We analyze empirical data to count the occurrences of each interaction type, and use this information as a signature to compute differences in higher-order connectivity patterns. Then, for each node, we investigate the overlap among its source and target sets, to extract recurring groups of co-senders and co-receivers. By examining this overlap and comparing it against randomized models, we aim to reveal whether certain systems exhibit a more redundant organization, where interactions frequently recur among the same groups, or a more diverse structure with less overlap among participants. Additionally, we propose new, computationally efficient definitions for reciprocity<sup>41</sup> for directed hypergraphs, namely exact, strong and weak higher-order

<sup>1</sup>Department of Network and Data Science, Central European University, Vienna, Austria. <sup>2</sup>Department of Information Engineering and Computer Science, University of Trento, Trento, Italy. ✉e-mail: [lotitoq@ceu.edu](mailto:lotitoq@ceu.edu); [battistonf@ceu.edu](mailto:battistonf@ceu.edu)

reciprocity, designed to capture different patterns of bi-directionality in empirical data. Finally, we extend motif analysis<sup>42</sup> to incorporate the directionality of interactions, extracting recurring higher-order and directed subgraphs. Our results suggest the existence of complex mechanisms of feedback and reinforcement in the information flow among system units, where pairwise interactions support the action of groups, and vice versa.

## Results

Traditional graph models reduce directed group interactions into a collection of pairwise links, often leading to a loss of important structural information about group organization and dynamics. For instance, reducing a many-to-many interaction, such as SOURCE = {A, B} and TARGET = {D, E} to a set of directed pairwise links ( $A \rightarrow D, A \rightarrow E, B \rightarrow D$  and  $B \rightarrow E$ ) fails to capture the collective nature of the interaction, including information about co-senders and co-receivers nodes. Directed hypergraphs preserve both group-based structure and the associated information flow, allowing for a more faithful representation of complex interactions. In such a framework, hyperedge direction is encoded by distinguishing between source and target node sets, which we consider non-empty and disjoint. More formally, we work with finite, simple, directed hypergraphs  $\mathcal{H} = (V, E, \mathbf{s}, \mathbf{t})$  on node set  $V$ , where each hyperedge  $e \in E$  is an ordered pair  $(\mathbf{s}(e), \mathbf{t}(e))$  with  $\mathbf{s}(e), \mathbf{t}(e) \subseteq V$ ,  $\mathbf{s}(e) \neq \emptyset$ ,  $\mathbf{t}(e) \neq \emptyset$  and  $\mathbf{s}(e) \cap \mathbf{t}(e) = \emptyset$ . This definition yields four canonical directed hyperedge patterns: one-to-one, where a single source node connects to a single target; one-to-many, where one source affects multiple targets; many-to-one, where multiple sources act on a single target; and many-to-many, the most general case, where multiple sources act on multiple targets. The analysis of the interplay and overlap among these building blocks in real-world hypergraphs enables a characterization of their microscale organization. Figure 1 illustrates this taxonomy on a toy directed hypergraph.

We analyzed datasets from multiple domains, including QNA (nodes are users and forum posts are hyperedges), E-MAIL (nodes are users and emails are hyperedges), BITCOIN (nodes are accounts and financial transactions are hyperedges), METABOLIC (nodes are genes and metabolic reactions are hyperedges) and CITATION (nodes are authors and hyperedges are paper citations)<sup>40</sup>. Each dataset is encoded as a set-indexed adjacency tensor. In particular, we index the distinct source- and target-sets observed in the data by  $\{S_\alpha\}_{\alpha=1}^p$  and  $\{T_\beta\}_{\beta=1}^q$ , respectively, and define the set-indexed adjacency tensor  $\mathcal{A} \in \{0, 1\}^{p \times q}$  by  $\mathcal{A}_{\alpha\beta} = 1$  if and only if there exists  $e \in E$  with  $\mathbf{s}(e) = S_\alpha$  and  $\mathbf{t}(e) = T_\beta$  (and  $\mathcal{A}_{\alpha\beta} = 0$  otherwise). Whenever  $\mathcal{A}_{\alpha\beta} = 1$ , we enforce  $S_\alpha \neq \emptyset$ ,  $T_\beta \neq \emptyset$  and  $S_\alpha \cap T_\beta = \emptyset$ .

Detailed descriptions and summary statistics of each dataset are reported in Supplementary Note 1.

### Patterns of directed hyperedges

A natural starting point to characterize directed hypergraphs across domains is investigating the diversity in their patterns of directed hyperedges. For each dataset, we construct a *hyperedge signature vector*, which captures the distribution of hyperedges based on the sizes of their source and target sets (see Methods). Such vectors provide a fingerprint for systems based on their higher-order connectivity patterns at the microscale. Figure 2a shows the hyperedge signature vectors for each dataset, considering interactions up to size 6. To emphasize the role of higher-order interactions in the analysis, we do not consider one-to-one interactions. We find that one-to-many interactions dominate the E-MAIL dataset, reflecting the typical structure of email communications. Similarly, in the QNA, many-to-one interactions are prevalent, as these systems involve multiple individuals responding to a question by a single user. In contrast, METABOLIC and CITATION datasets show high abundances in many-to-many relationships across a variety of source and target set sizes. Finally, BITCOIN dataset exhibits more varied behavior, with abundant entries for both one-to-many and many-to-many interactions, indicating different interaction types in the network.

To further explore structural diversity across different domains, we compute pairwise rank correlations among hyperedge signature vectors

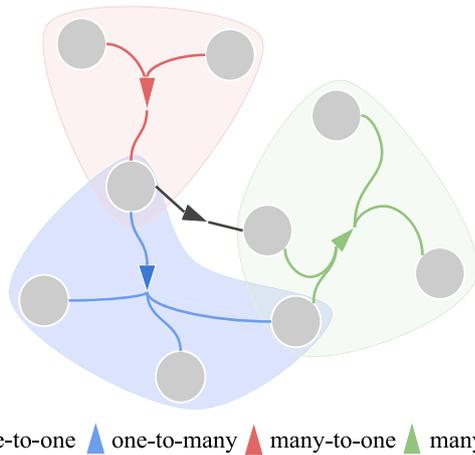
using weighted Kendall's  $\tau$  and apply hierarchical agglomerative clustering on their correlation matrix. A correlation value close to 1 indicates similar hyperedge structures, 0 suggests no relationship, and  $-1$  indicates the structures are inversely related. The clustering procedure applied to the systems' correlation matrix results in a dendrogram that visually represents their hierarchical relationships, highlighting the presence of clusters of directed hypergraphs that share similar connectivity patterns. In Fig. 2b, we show the correlation matrix and the clustering dendrogram. By examining the correlation matrix, we observe a strong correlation within systems from the same domain, indicating highly similar abundance in hyperedge structures. In contrast, systems from different domains exhibit varying degrees of correlation. Specifically, E-MAIL and QNA datasets are inversely correlated, as they display non-overlapping and complementary connectivity patterns: E-MAIL is characterized by one-to-many interactions, whereas QNA primarily involves many-to-one relationships. The METABOLIC and CITATION datasets, which feature many-to-many interactions, are positively correlated and form a distinct cluster. Interestingly, the BITCOIN datasets also display positive correlations with the METABOLIC and CITATION cluster due to a high presence of many-to-many interaction patterns. However, they also exhibit a weaker positive correlation with the E-MAIL datasets, reflecting the presence of one-to-many interactions in BITCOIN.

### Source and target sets overlap

The degree to which hyperedges share elements provides valuable insights into redundancy, hierarchical structures, and information flow within the system. In general, real-world systems exhibit a high degree of overlap, indicating that recurrent and redundant interactions are a shared feature<sup>43</sup>. Moreover, hyperedge overlap has been shown to widely impact the dynamics of systems with higher-order interactions<sup>44,45</sup>.

In directed hypergraphs, nodes are frequently involved in multiple hyperedges, either as part of the source or the target set. In order to characterize and quantify this property, for each node, we measure how much its incident source sets overlap and how much its incident target sets overlap, and we compare with a null model, reporting source and target z-scores (see Methods). In other words, we measure the extent to which nodes engage in interactions with the same set of co-senders or co-receivers. Specifically, an observed overlap associated with a z-score greater than 2 (i.e., statistically significant excess overlap) highlights nodes that tend to participate in structurally redundant interactions where the same groups of nodes frequently co-occur in source or target sets. Conversely, a z-score less than  $-2$  reflects statistically significant lower overlap than expected, suggesting that interactions are more diverse, with hyperedges being more distinct and less likely to share members.

In Fig. 3, we show the distribution of nodes with a given overlap z-score for source and target sets across domains, along with the fraction of nodes in each region of the z-score space. The focus is on positive excess overlap, as statistically significant negative excess overlap is very rare. The CITATION dataset displays significant overlap for both source and target sets, showing that (i) an author tends to preferentially work with known collaborators, and (ii) an author tends to be cited repeatedly alongside similar sets of authors. In the E-MAIL dataset, the excess overlap can be computed only for the target sets, as the source sets always have cardinality 1. The overlap is significantly larger than random, underscoring the hierarchical and broadcast-like nature of email communication. The BITCOIN dataset generally exhibits nodes with high excess overlap in both source and target sets. However, the presence of nodes with excess overlap values lower than zero implies that certain participants in the network engage in interactions that introduce more novelty rather than reinforce existing hyperedges. The METABOLIC dataset follows a similar trend, with half of the nodes displaying significant excess overlap for either source or target sets or both, suggesting that metabolic reactions tend to involve recurring sets of substrates and products and highlighting the modular nature of metabolic networks. QNA data show a lower excess overlap compared with the other systems, indicating that forum respondents are less likely to engage repeatedly with the same set



**Fig. 1 | Schematic of a directed hypergraph.** Each interaction encodes a source set of units acting towards a target set of units. We distinguish four types of directed higher-order interactions: one-to-one (black), one-to-many (blue), many-to-one (red), and many-to-many (green).

of co-responders. Since the target sets in this dataset always have cardinality 1, the excess overlap can be computed only for source sets.

**Higher-order reciprocity**

Reciprocity is a fundamental property of systems with directed interactions, including social networks<sup>46</sup>. It traditionally refers to the tendency of the system’s units to mutually exchange information. In directed graphs, reciprocity is defined as

$$r = \frac{\tilde{L}}{L},$$

measuring the ratio of the number of bidirectional links ( $\tilde{L}$ ) to the total number of links ( $L$ ). This measure is often normalized as

$$\rho_{NM} = \frac{r - \langle r \rangle_{NM}}{1 - \langle r \rangle_{NM}}$$

where  $r$  is the observed reciprocity in the network, and  $\langle r \rangle_{NM}$  is the average reciprocity in null model samples<sup>41</sup>. It measures the difference between observed and expected reciprocity by the maximum possible deviation, bounding  $\rho_{NM}$  between  $-1$  and  $1$ . Positive values indicate more reciprocity than expected at random, negative values indicate less, and values near zero suggest consistency with the null model. This normalization allows a more faithful comparison and ranking of reciprocity across systems with different scales and density<sup>41</sup>. Recognizing its broad importance, recent works have extended reciprocity to hypergraphs, accounting for the complexity of having multiple nodes in both the source and target sets of hyperedges. Among the recent approaches for hypergraph reciprocity, one method decomposes hyperedges into pairwise links<sup>39</sup>, losing information about group interactions. An alternative approach defines a more complex measure that diverges from the traditional binary definition of reciprocity at the level of single links<sup>40</sup>. While this approach can capture different nuances, it is computationally expensive and less straightforward to interpret, as it provides a continuous value instead of a simple yes-or-no answer to whether an interaction is reciprocated.

Here, we introduce three simple and computationally efficient measures for higher-order reciprocity in directed hypergraphs, capturing different aspects of mutual interactions:

- Exact reciprocity occurs when an interaction represented by a hyperedge with a source set  $h$  and a target set  $t$  is precisely mirrored by another interaction with the source and target sets reversed. Formally,

two hyperedges  $e_1 = (h_1, t_1)$  and  $e_2 = (h_2, t_2)$  are exactly reciprocated if and only if  $h_1 = t_2$  and  $t_1 = h_2$ . This is the strictest form of reciprocity.

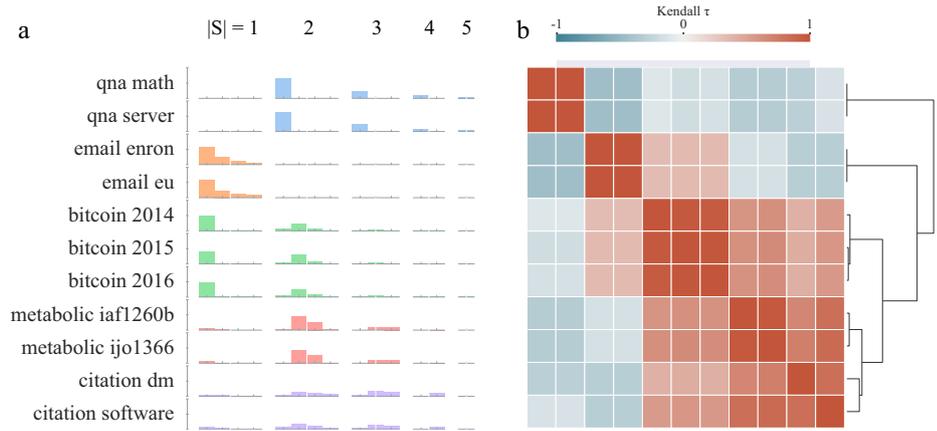
- Strong reciprocity relaxes the previous requirement and allows source and target sets to be reversed through a combination of hyperedges, instead of requiring a direct reversal with a single opposite one. Formally, a hyperedge  $e = (h, t)$  is *strongly reciprocated* if there exists a set of hyperedges  $\{e_1, e_2, \dots, e_k\}$  such that the union of the target sets of  $e_1, \dots, e_k$  is a superset of the source set  $h$ , and the union of the source sets of  $e_1, \dots, e_k$  is a superset of the target set  $t$ .
- Weak reciprocity represents the most relaxed form of reciprocity and requires only that at least one node from the target set of a hyperedge appears in the source set of another, and vice versa. Formally, a hyperedge  $e = (h, t)$  is *weakly reciprocated* if there exists another hyperedge  $e' = (h', t')$  such that  $h \cap t' \neq \emptyset$  and  $t \cap h' \neq \emptyset$ .

We summarize our definitions of reciprocity for directed hypergraphs in Fig. 4. More information about the algorithmic aspects of such measures is available in the Methods section.

After introducing these definitions, a natural first question is which systems exhibit the highest and lowest levels of reciprocity and how the ranking of systems based on reciprocity changes across different definitions. We address this in Fig. 5a, which shows the normalized ratio of reciprocated hyperedges  $\rho_{NM}$  (reciprocity score) for each system across varying notions of reciprocity. The reciprocity score induces rankings of the systems, allowing us to observe which systems exhibit stronger tendencies toward mutual exchange of information. By definition, the score tends to increase for each system as we move from stricter definitions of reciprocity (exact) to more relaxed ones (weak). We observe that systems from the same domain tend to show similar levels of reciprocity across definitions, indicating that functional similarities within domains may drive comparable reciprocity patterns. E-MAIL datasets exhibit the highest levels of reciprocity, while BITCOIN datasets consistently show the lowest. Interestingly, while the ranking of systems remains largely stable with varying definitions, the relative distances between the datasets change. For instance, exact reciprocity mostly characterizes E-MAIL datasets, which are positioned far from the other datasets, clustering distinctly at the top of the scale. Strong reciprocity induces three clear clusters of datasets based on their scores: E-MAIL datasets rank the highest by a large margin, while BITCOIN datasets occupy the very low end. In the case of weak reciprocity, the datasets begin to separate along domain lines, spanning the entire spectrum of reciprocity scores. Notably, we observe a reduction in the distance between E-MAIL and datasets from metabolic and citation domains, suggesting a convergence in reciprocity levels as the definition becomes more relaxed. Overall, these patterns highlight how the choice of measure can influence the perceived level of reciprocity within different systems. By analyzing how the score evolves across definitions, we gain a more precise understanding of the extent of mutual exchange within each system, from the high reciprocity observed in E-MAIL datasets, where high mutual exchange is clear, to the lower reciprocity in BITCOIN datasets, where reciprocal connections are minimal across all definitions, and to the metabolic datasets, which emerge with high reciprocity under weaker definitions.

A related question is how the size of hyperedges influences the levels of reciprocity. We address this in Fig. 5b, which reports the reciprocity score as a function of interaction size for the three notions, with deviations interpreted relative to the null baseline. For the E-MAIL datasets,  $\rho_{NM}$  displays consistent positive offset for strong and weak reciprocity across all admissible sizes, whereas exact reciprocity exhibits a pronounced signal only at size 2 and is negligible at larger sizes. This suggests that excess reciprocal structure relative to the null is present for most interaction sizes, while exact “mirror” reversals remain largely a dyadic feature. In the METABOLIC datasets, exact reciprocity remains near zero across sizes, with strong and weak deviations showing initial negative values for small sizes, followed by a substantial increase at higher sizes that is more pronounced for the weak definition. Thus, small-set interactions are under-reciprocated relative to the null, whereas large-set interactions become increasingly over-

**Fig. 2 | Hyperedge signature of directed hypergraphs. a** We describe each system with a hyperedge signature vector whose entries encode the count of directed hyperedges with source and tail sizes ( $|S|, |T|$ ). We compute statistics using hyperedges with total cardinality at most 6 (i.e.,  $|S| + |T| \leq 6$ ). For visualization, we display each vector as a sequence of histogram panels: one panel for each source size  $|S| = i$ , separated by a small gap; within each panel, bins correspond to tail sizes  $|T|$  in increasing order, restricted by  $i + |T| \leq 6$ . Systems from the same domain share the color. **b** Dendrogram resulting from agglomerative clustering applied to the correlation matrix of hyperedge signature vectors for each dataset. Correlation values are color-coded, with high positive correlations in red and high negative correlations in blue.



represented. In the CITATION datasets, exact reciprocity is modestly positive at small sizes. Strong reciprocity displays a higher, stable excess across sizes, while weak reciprocity reaches its highest value for the smallest sets and diminishes with size. Citation reciprocity above the null is thus dominated by weak reciprocity in small interactions, with strong reciprocity contributing a more size-invariant offset. For the BITCOIN datasets, scores for all definitions cluster near zero, with a slight increase from exact to strong to weak, and minimal size dependence. This points to weak, size-stable excess reciprocal structure. Finally, the QNA datasets track the BITCOIN pattern, with modest positive offsets that increase from exact to strong to weak definitions, and little systematic variation as size increases.

These findings suggest that weaker notions of reciprocity are valuable in providing insights into the overall reciprocity of systems with larger interactions, capturing a multifaceted view of how reciprocity operates at different strengths and interaction sizes.

### Motif analysis in directed hypergraphs

Motif analysis involves counting the frequency of patterns of interactions in connected subgraphs of a given number of nodes. This framework was first introduced by Milo et al.<sup>42</sup> to extract the fundamental functional units of complex systems<sup>47</sup>. Recently, motif analysis has been extended to undirected hypergraphs to capture patterns of interactions with arbitrary size<sup>20</sup>. Here, we extend such analysis to consider also the direction of the hyperedges involved in the patterns.

First, it is interesting to study the combinatorics of the patterns of directed subhypergraphs. There is no simple closed-form formula for counting the number of possible directed higher-order motifs as a function of their order  $n$ , i.e., the number of nodes in the patterns. We can estimate the number of non-isomorphic connected directed hypergraphs in a way similar to ref. 20. Given a set of  $n$  nodes, the number of possible directed hyperedges is  $3^n - 2 \cdot \sum_{k=1}^n \binom{n}{k} - 1 = 3^n - 2 \cdot 2^n + 1$ . This expression counts the ways to partition the  $n$  nodes into three disjoint sets: source, target and empty set. We subtract the invalid combinations with empty source or target sets. Given  $n$  nodes, we ensure connectivity by selecting a chain of  $n - 1$  hyperedges and including them in the hypergraph, leaving us with  $3^n - 2 \cdot 2^n - n + 2$  remaining possible hyperedges. For each remaining hyperedge, we decide whether to include it or not, resulting in  $2^{3^n - 2 \cdot 2^n - n + 2}$  total hypergraphs. Since we are interested in non-isomorphic hypergraphs, we divide this number by  $n!$ , the number of ways to label the vertices, providing the lower bound  $\frac{2^{3^n - 2 \cdot 2^n - n + 2}}{n!}$ . If we ignore the constraints of non-isomorphism and connectivity, we count the number of possible labeled hypergraphs. Since each of the  $3^n - 2 \cdot 2^n + 1$  possible hyperedges can either be included or excluded, the total number of labeled hypergraphs is at most  $2^{3^n - 2 \cdot 2^n + 1}$ . Figure 6 shows the upper and lower bounds on the growth of possible sub-hypergraph patterns as a function of the number of nodes

(order) for both the undirected and directed cases. The estimated number of patterns grows super-exponentially, even in the undirected case. In the directed case, the growth is even faster due to the need to consider all possible subdivisions into source and target sets.

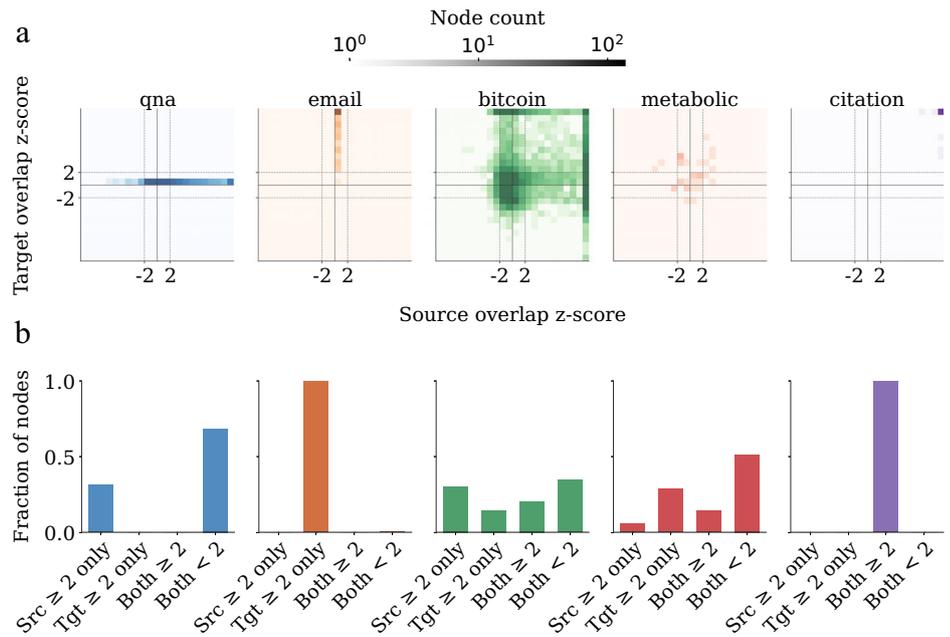
To perform motif analysis on real-world directed hypergraphs, we propose an exact algorithm to count the frequency of all connected sub-hypergraph patterns and associate each pattern with a z-score that quantifies its over- or under-representation relative to our proposed configuration model (see Methods).

Given the intractability of the problem for large sub-hypergraphs, we limit our study of empirical data to patterns involving three and four nodes. Moreover, we focus on patterns that include at least one group interaction. Because the theoretical pattern space is still very large, while the empirical signal is sparse and concentrated (see Supplementary Note 3), we restrict our analysis to the positive tail of the motif z-score distributions (see Supplementary Note 3) to highlight structural differences across datasets. In particular, Fig. 7 shows the most over-represented patterns of directed higher-order interactions with three and four nodes across different domains. Each domain reveals distinct motifs, characterized by different directed hyperedge types, sizes, densities and patterns of reciprocity. In terms of hyperedges types, E-MAIL and QNA involve abundant patterns with only many-to-one and one-to-many interactions. Other datasets display more diverse patterns, including combinations of one-to-many, many-to-one, and many-to-many interactions (the latter is possible only in motifs with four nodes). Traditional one-to-one interactions are commonly part of abundant patterns in all datasets. The number of interactions in abundant sub-hypergraphs is small in the BITCOIN, METABOLIC and CITATION domains, often involving just one or two hyperedges. In contrast, the E-MAIL and QNA domains tend to be richer in interactions. This observation is reversed when considering the average size of interactions. The relation between the number and the average size of interactions aligns with previous studies on undirected higher-order motifs<sup>20</sup>. A common pattern in many datasets is the coexistence of group interactions alongside lower-order interactions within the same set of nodes, a phenomenon often referred to as nestedness<sup>20</sup> or simpliciality<sup>48</sup>. Interestingly, considering the direction of such interactions reveals that they seem to play a role in increasing the overall reciprocity of the patterns, suggesting the existence of a feedback mechanism. This is particularly evident in E-MAIL data. In addition to reciprocity, the direction of lower-order interactions in abundant patterns suggests a reinforcing mechanism where subsets of source and target nodes interact at multiple interaction sizes. These observations are closely connected with the insights discussed in the previous sections about frequent co-senders and co-receivers nodes and higher-order reciprocity.

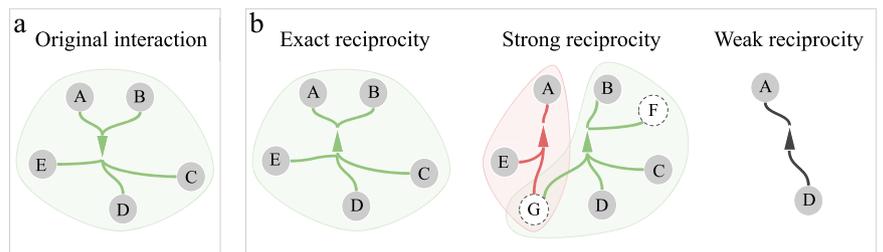
### Discussion

Hypergraphs extend traditional network representations by allowing hyperedges to connect multiple nodes simultaneously, enabling the

**Fig. 3 | Overlap across domains.** **a** Distribution of node counts within the joint z-score space of source and target overlap, representing how much nodes deviate from null model expectations in both dimensions. **b** Bar plots quantifying the fraction of nodes exceeding the z-score  $\geq 2$  threshold for either source or target overlaps, nodes exceeding both thresholds, and nodes below both thresholds.



**Fig. 4 | Reciprocity measures for directed hypergraphs.** **a** Example of a directed hyperedge. **b** All possible ways in which this hyperedge can be reciprocated according to our definitions. Exact reciprocity: a single hyperedge with source and target sets swapped, represented by reversing the arrow between the same node sets. Strong reciprocity: multiple hyperedges collectively reverse the interaction, possibly involving external nodes. Weak reciprocity: at least one node in the target set reciprocates with one node in the source set, illustrated as a pairwise link with reversed direction. In all panels, shaded areas group nodes involved in each interaction; colors encode the interaction pattern (green many-to-many, red many-to-one, black one-to-one); arrows encode direction from source to target. Gray disks denote nodes in the original hyperedge, and external nodes that appear only in reciprocal interactions are white with a dashed border.



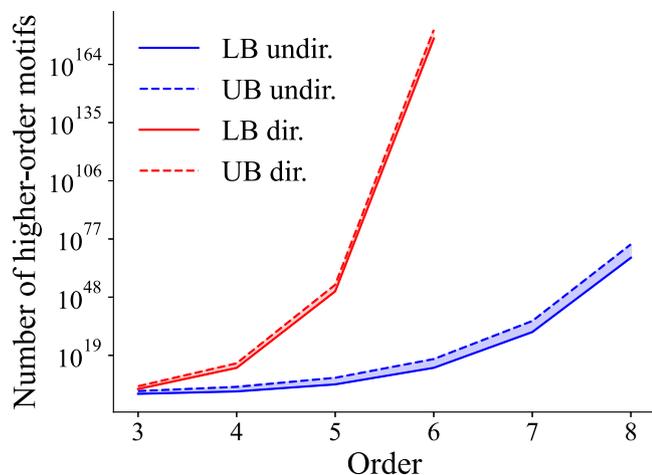
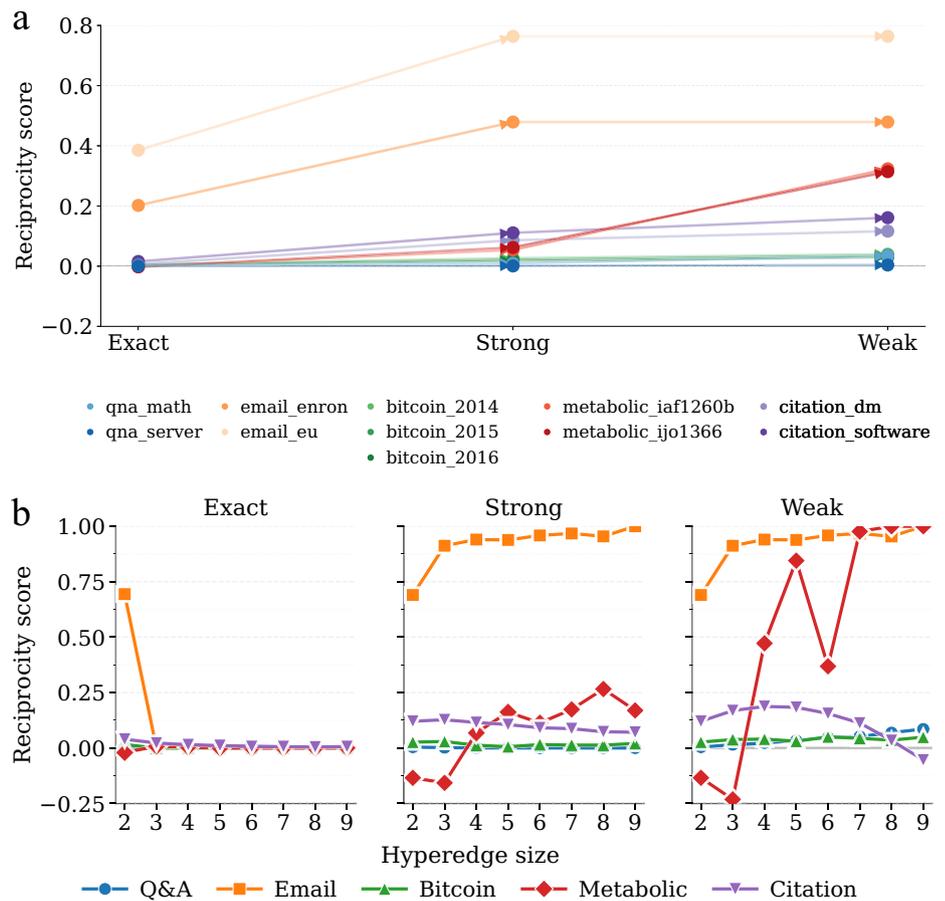
encoding of group interactions ubiquitous in many relational systems. Directed hypergraphs further enhance our modeling abilities by accounting for directionality in group interactions, distinguishing between source and target sets for each hyperedge. This versatile framework can accurately model a range of diverse real-world systems and interactions, including financial transactions, email exchanges, and metabolic reactions.

In this work, we proposed new measures and tools to analyze the structural organization of directed hypergraphs at their microscale. First, we analyzed hyperedge signature vectors to identify the abundance of each hyperedge structure across datasets and identified classes of systems sharing similar higher-order connectivity patterns. Second, we analyzed the excess overlap among source and target sets for each node in each system. The resulting distributions suggest that different domains may follow distinct organizational principles, ranging from redundant to more diverse interaction patterns. Then, we introduced three distinct types of higher-order reciprocity measures: exact, strong, and weak reciprocity. Each definition offers a different perspective on how group interactions can be reciprocated, ranging from strict to more relaxed forms of reciprocal influence, and can be computed efficiently, making it suitable also for the analysis of very large

systems. We showed that all systems exhibit reciprocity in broad terms, though different domains are associated with specific patterns and sensitivity to specific reciprocity measures. Lastly, we extended the notion of motifs to directed hypergraphs, capturing recurring patterns of directed interactions. Motif analysis revealed frequent microscale structures and highlighted common organizational principles playing a role in the function and behavior of systems, such as the existence of reinforcing or feedback mechanisms among dyadic and non-dyadic interactions in groups.

Taken together, by considering the nuances related to the directionality of interactions in directed hypergraphs, our research provides a framework to understand higher-order connectivity in directed complex systems, opening up a wide range of potential applications in diverse fields, such as social network analysis, biology, and finance. For instance, the study of multi-party financial transactions as directed higher-order structures may capture more complex patterns of fraudulent activity than traditional graph-based models<sup>49</sup>. Similarly, directed hypergraphs may enhance the accuracy of existing frameworks in identifying and predicting important genes based on genomic expression relations<sup>50</sup>. As scalability is a pressing issue in hypergraph algorithms, future work may explore more advanced

**Fig. 5 | Higher-order reciprocity in real-world hypergraphs.** **a** Reciprocity score across datasets and reciprocity definitions. Each column corresponds to a distinct notion of higher-order reciprocity, thereby inducing a ranking of the datasets based on their scores. Datasets from the same domains share the same color. Arrows link the datasets across different definitions. **b** Reciprocity score disaggregated by hyperedge size for each different notion of reciprocity. Trends should be interpreted relative to the null (zero line). To simplify the plots, we aggregate systems from the same domain.



**Fig. 6 | Combinatorics of directed higher-order motifs.** Upper (dashed lines) and lower (solid lines) bounds on the number of higher-order motifs as a function of their order. Blue lines refer to undirected motifs on hypergraphs, red lines refer to the directed case.

techniques for detecting motifs in large-scale directed hypergraphs. These include sampling methods already proposed for undirected higher-order motifs<sup>51</sup> or the use of parallel algorithms, which may achieve significant speed-ups beyond the implementation of our current Python library, thereby enabling analysis of motifs larger than four nodes. Another interesting venue for further studies is related to the study of reciprocity in time-

evolving hypergraphs, since it can affect mechanisms of group formation<sup>4,5</sup> and inform the efficient seeding of information<sup>52,53</sup>. All in all, our work reveals new structural principles behind the organization of real-world systems, shedding light on the complex interplay between structural patterns and functionality in directed complex systems.

## Methods

### Hyperedge signature vector construction

For each dataset, we construct a *hyperedge signature vector*  $\mathbf{v}$ , where each element represents the count of hyperedges with a specific combination of source set size  $s$  and target set size  $t$  in the hypergraph. The vector  $\mathbf{v}$  captures the distribution of hyperedges based on the sizes of their source and target sets, providing a profile of the hypergraph structure.

Formally, we define the vector  $\mathbf{v}$  as follows:

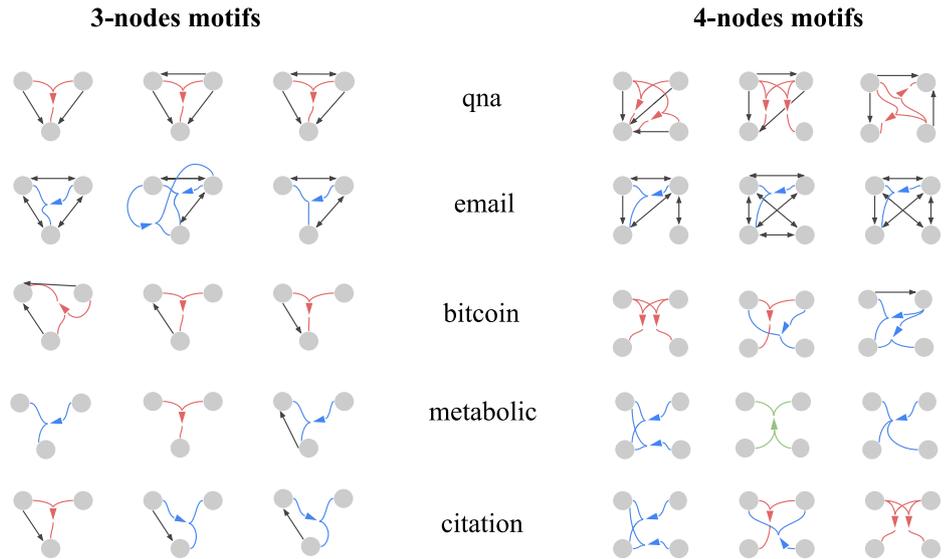
$$\mathbf{v} = (v_{1,2}, \dots, v_{1,K-1}, v_{2,1}, \dots, v_{2,K-2}, \dots, v_{K-1,1})$$

where  $K$  represents the maximum hyperedge size considered, and each  $v_{h,t}$  counts the number of hyperedges with a specific source size  $h$  and target size  $t$ .

### Microcanonical set-swap configuration model for directed hypergraphs

We generate randomized counterparts of directed hypergraphs using a microcanonical set-swap configuration model. This approach is similar to the one proposed by Preti et al.<sup>36</sup>, and more recently by Kraakman et al.<sup>54</sup>. More broadly, our null model fits within the family of configuration- and entropy-based random hypergraph ensembles that preserve degree and edge-size sequences<sup>55-58</sup>, which have been used in applications, such as

**Fig. 7 | Directed higher-order motifs in real-world hypergraphs.** The three most representative directed higher-order motifs of orders three and four from each system. The color of a group interaction encodes its type: one-to-one (black), one-to-many (blue), many-to-one (red), and many-to-many (green). We group statistics of systems within the same domain.



extracting statistically validated higher-order interactions<sup>59</sup> and motifs<sup>20</sup>. In our framework, each directed hyperedge  $e \in E$  is specified by two disjoint node sets: a source set  $s(e) \subseteq V$  and a target set  $t(e) \subseteq V$ , with  $s(e) \cap t(e) = \emptyset$ . The microcanonical ensemble comprises all simple directed hypergraphs that match the observed node out-/in-degree sequences (counts of appearances in  $s(\cdot)$  and  $t(\cdot)$ ), respectively) and preserve every hyperedge's cardinalities  $|s(e)|$  and  $|t(e)|$ . We forbid duplicated hyperedges.

Starting from the observed  $H$ , we repeatedly attempt *set-swaps* on one side at a time. In a source-side attempt, we select two distinct hyperedges  $e_i \neq e_j$  uniformly and sample  $u \in s(e_i)$  and  $v \in s(e_j)$ . We propose

$$s'(e_i) = (s(e_i) \setminus \{u\}) \cup \{v\}, \quad s'(e_j) = (s(e_j) \setminus \{v\}) \cup \{u\},$$

leaving  $t(e_i)$  and  $t(e_j)$  unchanged. The move is accepted only if it respects set semantics on the chosen side (no duplicates), preserves within-edge disjointness ( $s(e) \cap t(e) = \emptyset$ ), and does not create a duplicate ( $s'(e), t(e)$ ). Target-side attempts are defined symmetrically. Accepted swaps conserve each node's in/out degree and all  $|s(e)|, |t(e)|$ . Rejected proposals leave the state unchanged.

**Hyperedge overlap**

For each node, we quantify its hyperedge overlap separately for its participation in target sets (in-hyperedges) and source sets (out-hyperedges). For in-hyperedge overlap, we consider all hyperedges  $e$  for which the node is in the target set  $t(e)$ . Let  $\mathcal{E}_{in}$  be the collection of these hyperedges. We define the in-hyperedge overlap as

$$O_{in} = \frac{\sum_{e \in \mathcal{E}_{in}} |t(e)|}{|\mathcal{E}_{in}| |\bigcup_{e \in \mathcal{E}_{in}} t(e)|},$$

where  $|t(e)|$  denotes the number of nodes in the target set of hyperedge  $e$ .

Similarly, for out-hyperedge overlap, we consider only hyperedges  $e$  in which the node is a source (i.e., belongs to  $s(e)$ ). Let  $\mathcal{E}_{out}$  denote this collection, and define

$$O_{out} = \frac{\sum_{e \in \mathcal{E}_{out}} |s(e)|}{|\mathcal{E}_{out}| |\bigcup_{e \in \mathcal{E}_{out}} s(e)|},$$

with  $|s(e)|$  being the size of the source set of hyperedge  $e$ . These metrics yield a value of 1 when all corresponding hyperedges share an identical set of nodes (i.e., maximal overlap), and decrease as the sets become more diverse.

To assess statistical significance, we compare the observed overlaps to those computed on an ensemble of randomized networks that preserve key structural properties (e.g., node degrees and hyperedge sizes). For each node  $i$ , we compute a z-score that standardizes the observed overlap relative to the null ensemble:

$$z_O(i) = \frac{O_i^{obs} - \langle O_i^{rand} \rangle}{\sqrt{\text{Var}(O_i^{rand})}}.$$

We compute this for both  $O_{sources}$  and  $O_{targets}$ .

**Algorithms for measuring reciprocity in directed hypergraphs**

Below, we outline our proposed algorithms for efficiently measuring reciprocity in directed hypergraphs.

- **Exact reciprocity.** Each hyperedge  $e = (s, t)$  is stored in a hash-based dictionary, and for each hyperedge, we search for a reverse hyperedge  $e' = (t, s)$ . Since each lookup takes constant time, the overall complexity is  $O(m)$ , where  $m$  is the number of hyperedges.
- **Strong reciprocity.** For each hyperedge  $e = (s, t)$ , we maintain a reachability dictionary that tracks which nodes in the target set  $t$  can reach other nodes via multiple hyperedges. We then check whether the source set  $s$  is fully covered by the accumulated reachable nodes from the target set  $t$ . This involves iterating over each hyperedge, for each target node, accumulating the reachable nodes and then checking if the source set is a subset of this accumulated set. Computing the union of reachable nodes is  $O(s \cdot t)$ , where  $s$  is the maximum size of source sets and  $t$  is the maximum size of target sets. This operation is repeated for all hyperedges, leading to a total complexity of  $O(m \cdot s \cdot t)$ .
- **Weak reciprocity.** First, we construct a dictionary to store all directed node pairs between the source and target sets of each hyperedge. Then, for each hyperedge, we check whether any of its target nodes are linked back to the source nodes via reverse connections in the dictionary. The computational complexity is dominated by the first operation, which is  $O(m \cdot s \cdot t)$ , where  $s$  is the maximum size of the source sets and  $t$  is the maximum size of the target sets across all hyperedges.

In practice, executing these algorithms on the real-world datasets used in our experiments requires only a few minutes for all datasets combined, demonstrating the computational efficiency of the proposed methods.

### Algorithms for motif analysis in directed hypergraphs

In order to design efficient algorithms for mining directed higher-order motifs, we extend prior ideas developed for the same problem in undirected hypergraphs<sup>51</sup>. Our algorithms are efficient enough to count motifs of size 3 and 4 in datasets of reasonable size (comparable to those used in our experiments). However, scaling to larger datasets and motifs of larger size would require more sophisticated approaches, such as sampling algorithms<sup>51</sup>, which we leave for future work. Further details on the execution times of the algorithms for mining motifs of orders 3 and 4 can be found in Supplementary Notes 2.

The algorithm for mining motifs (involving at least one group interaction) of order 3 begins by iterating through each hyperedge in the hypergraph that contains exactly three vertices. For each such hyperedge, it identifies all possible subsets of vertices and checks whether one or more subsets form valid directed hyperedges in the hypergraph. Valid subsets, along with the original hyperedge, define the motif structure involving those three vertices. To ensure consistency in motif identification, the algorithm generates a canonical form of the motif by lexicographical ordering its vertices and edges, which can be computed by sorting the  $n!$  possible re-labels. This canonical representation allows motifs with the same structural pattern to be compared and counted, even if they differ in their vertex labels. Each canonical form of motifs is stored in a frequency hash map. If the motif has not been encountered before, it is added to the map; if it has, its frequency count is incremented. In the end, the algorithm outputs a distribution of the various motif structures of order 3. This algorithm operates in linear time with respect to the number of hyperedges of order 3. Specifically, its computational complexity is  $O(m_3)$ , where  $m_3$  is the number of hyperedges involving exactly three vertices. Each motif construction and comparison is performed in constant time due to the fixed size of the motifs. For more details, refer to the pseudocode in Supplementary Note 2. The algorithm for mining motifs of order 4 follows a similar approach. First, it iterates over all hyperedges of size 4, counting the motifs involving exactly these 4 nodes. Unlike the previous algorithm, it then iterates over all hyperedges of size 3, performing an additional neighborhood exploration step to identify the fourth node involved in the motif. Each neighboring node is considered during this process. Once the 4 nodes are identified, the algorithm constructs the motif as before. The pseudocode for this algorithm is provided in Supplementary Note 2.

### Statistical significance of motifs

To distinguish meaningful, non-random interaction patterns from those that may occur by chance, we use a configuration model as a null model to evaluate the statistical significance of the interaction patterns after computing their frequency in our directed hypergraphs. The configuration model generates randomized versions of the original hypergraph while preserving key properties, such as the in-degree and out-degree sequences, as well as the source and target sizes of the hyperedges<sup>36</sup>. By comparing the observed frequencies with those found in the randomized networks, we can identify significantly over-represented motifs.

In particular, each motif  $i$  is associated with a standardized score  $z_M(i)$ , which quantifies how many standard deviations the observed frequency of the motif differs from its expected value under the configuration model<sup>20,42</sup>. A larger absolute value of  $z_M(i)$  indicates a stronger statistical deviation from the null expectation, meaning the motif is significantly over- or under-represented compared to random networks. The score is defined as:

$$z_M(i) = \frac{N_i^{\text{obs}} - \langle N_i^{\text{rand}} \rangle}{\sqrt{\text{Var}(N_i^{\text{rand}})}}.$$

Here,  $N_i^{\text{obs}}$  denotes the observed frequency of motif  $i$  in the empirical hypergraph, and  $\langle N_i^{\text{rand}} \rangle$  and  $\text{Var}(N_i^{\text{rand}})$  are the mean and variance of motif  $i$ 's frequency across random hypergraphs generated by the configuration model. Following ref. 47, we generate 100 random samples from the

configuration model for each hypergraph to estimate the mean and variance of the motif frequencies.

### Data availability

Data<sup>40</sup> is publicly available and also easily accessible through HGX<sup>60</sup>.

### Code availability

The tools for the analysis of directed hypergraphs presented in this work are available as part of Hypergraphx (HGX)<sup>60</sup>.

Received: 27 March 2025; Accepted: 17 December 2025;

Published online: 07 January 2026

### References

- Boccaletti, S., Latora, V., Moreno, Y., Chavez, M. & Hwang, D.-U. Complex networks: Structure and dynamics. *Phys. Rep.* **424**, 175–308 (2006).
- Cimini, G. et al. The statistical physics of real-world networks. *Nat. Rev. Phys.* **1**, 58–71 (2019).
- Patania, A., Petri, G. & Vaccarino, F. The shape of collaborations. *EPJ Data Sci.* **6**, 1–16 (2017).
- Cencetti, G., Battiston, F., Lepri, B. & Karsai, M. Temporal properties of higher-order interactions in social networks. *Sci. Rep.* **11**, 1–10 (2021).
- Iacopini, I., Karsai, M. & Barrat, A. The temporal dynamics of group interactions in higher-order social networks. *Nat. Commun.* **15**, 7391 (2024).
- Ghoshal, G., Zlatić, V. & Caldarelli, G. Random hypergraphs and their applications. *Phys. Rev. E* **79**, 066118 (2009).
- Grilli, J., Barabás, G., Michalska-Smith, M. J. & Allesina, S. Higher-order interactions stabilize dynamics in competitive network models. *Nature* **548**, 210–213 (2017).
- Jost, J. & Mulas, R. Hypergraph laplace operators for chemical reaction networks. *Adv. Math.* **351**, 870–896 (2019).
- Traversa, P., Ferraz de Arruda, G., Vazquez, A. & Moreno, Y. Robustness and complexity of directed and weighted metabolic hypergraphs. *Entropy* **25**, 1537 (2023).
- Petri, G. et al. Homological scaffolds of brain functional networks. *J. R. Soc. Interface* **11**, 20140873 (2014).
- Santoro, A., Battiston, F., Lucas, M., Petri, G. & Amico, E. Higher-order connectomics of human brain function reveals local topological signatures of task decoding, individual identification, and behavior. *bioRxiv* 2023–12 (2023).
- Berge, C. *Graphs and hypergraphs* (North-Holland Pub. Co., 1973).
- Battiston, F. et al. Networks beyond pairwise interactions: structure and dynamics. *Phys. Rep.* **874**, 1–92 (2020).
- Battiston, F. et al. The physics of higher-order interactions in complex systems. *Nat. Phys.* **17**, 1093–1098 (2021).
- Benson, A. R. Three hypergraph eigenvector centralities. *SIAM J. Math. Data Sci.* **1**, 293–312 (2019).
- Tudisco, F. & Higham, D. J. Node and edge nonlinear eigenvector centrality for hypergraphs. *Commun. Phys.* **4**, 1–10 (2021).
- Eriksson, A., Edler, D., Rojas, A., de Domenico, M. & Rosvall, M. How choosing random-walk model and network representation matters for flow-based community detection in hypergraphs. *Commun. Phys.* **4**, 1–12 (2021).
- Contisciani, M., Battiston, F. & De Bacco, C. Inference of hyperedges and overlapping communities in hypergraphs. *Nat. Commun.* **13**, 1–10 (2022).
- Ruggeri, N., Contisciani, M., Battiston, F. & De Bacco, C. Community detection in large hypergraphs. *Sci. Adv.* **9**, eadg9159 (2023).
- Lotito, Q. F., Musciotto, F., Montresor, A. & Battiston, F. Higher-order motif analysis in hypergraphs. *Commun. Phys.* **5**, 79 (2022).
- Lee, G., Ko, J. & Shin, K. Hypergraph motifs: concepts, algorithms, and discoveries. *Proc. VLDB Endow.* **13**, 2256–2269 (2020).

22. Arregui-García, B., Longa, A., Lotito, Q. F., Meloni, S. & Cencetti, G. Patterns in temporal networks with higher-order egocentric structures. *Entropy* **26**, 256 (2024).
23. Petri, G. & Barrat, A. Simplicial activity driven model. *Phys. Rev. Lett.* **121**, 228301 (2018).
24. Di Gaetano, L., Battiston, F. & Starnini, M. Percolation and topological properties of temporal higher-order networks. *Phys. Rev. Lett.* **132**, 037401 (2024).
25. Gallo, L., Lacasa, L., Latora, V. & Battiston, F. Higher-order correlations reveal complex memory in temporal hypergraphs. *Nat. Commun.* **15**, 4754 (2024).
26. Schaub, M. T., Benson, A. R., Horn, P., Lippner, G. & Jadbabaie, A. Random walks on simplicial complexes and the normalized hodge 1-laplacian. *SIAM Rev.* **62**, 353–391 (2020).
27. Carletti, T., Battiston, F., Cencetti, G. & Fanelli, D. Random walks on hypergraphs. *Phys. Rev. E* **101**, 022308 (2020).
28. Lucas, M., Cencetti, G. & Battiston, F. Multiorder laplacian for synchronization in higher-order networks. *Phys. Rev. Res.* **2**, 033410 (2020).
29. Gambuzza, L. V. et al. Stability of synchronization in simplicial complexes. *Nat. Commun.* **12**, 1–13 (2021).
30. Zhang, Y., Lucas, M. & Battiston, F. Higher-order interactions shape collective dynamics differently in hypergraphs and simplicial complexes. *Nat. Commun.* **14**, 1605 (2023).
31. Iacopini, I., Petri, G., Barrat, A. & Latora, V. Simplicial models of social contagion. *Nat. Commun.* **10**, 1–9 (2019).
32. Chowdhary, S., Kumar, A., Cencetti, G., Iacopini, I. & Battiston, F. Simplicial contagion in temporal higher-order networks. *J. Phys.: Complex.* **2**, 035019 (2021).
33. Civilini, A., Sadekar, O., Battiston, F., Gómez-Gardeñes, J. & Latora, V. Explosive cooperation in social dilemmas on higher-order networks. *Phys. Rev. Lett.* **132**, 167401 (2024).
34. Ranshous, S. et al. Exchange pattern mining in the bitcoin transaction directed hypergraph. In *Financial Cryptography and Data Security: FC 2017 International Workshops, WAHC, BITCOIN, VOTING, WTSC, and TA, Sliema, Malta, April 7, 2017, Revised Selected Papers 21*, 248–263 (Springer, 2017).
35. Gallo, G., Longo, G., Pallottino, S. & Nguyen, S. Directed hypergraphs and applications. *Discret. Appl. Math.* **42**, 177–201 (1993).
36. Preti, G., Fazzino, A., Petri, G. & De Francisci Morales, G. Higher-order null models as a lens for social systems. *Phys. Rev. X* **14**, 031032 (2024).
37. Gallo, L. et al. Synchronization induced by directed higher-order interactions. *Commun. Phys.* **5**, 263 (2022).
38. Moon, H., Kim, H., Kim, S. & Shin, K. Four-set hypergraphlets for characterization of directed hypergraphs. *arXiv preprint arXiv:2311.14289* (2023).
39. Pearcy, N., Crofts, J. J. & Chuzhanova, N. Hypergraph models of metabolism. *Int. J. Biol. Vet. Agric. Food Eng.* **8**, 752–756 (2014).
40. Kim, S., Choe, M., Yoo, J. & Shin, K. Reciprocity in directed hypergraphs: measures, findings, and generators. *Data Min. Knowl. Discov.* **37**, 2330–2388 (2023).
41. Garlaschelli, D. & Loffredo, M. I. Patterns of link reciprocity in directed networks. *Phys. Rev. Lett.* **93**, 268701 (2004).
42. Milo, R. et al. Network motifs: Simple building blocks of complex networks. *Science* **298**, 824–827 (2002).
43. Lee, G., Choe, M. & Shin, K. How do hyperedges overlap in real-world hypergraphs?—patterns, measures, and generators. In *Proc. of the web conference*, 3396–3407 (ACM, 2021).
44. Malizia, F., Lamata-Otín, S., Frasca, M., Latora, V. & Gómez-Gardeñes, J. Hyperedge overlap drives explosive transitions in systems with higher-order interactions. *Nat. Commun.* **16**, 555 (2025).
45. Lamata-Otín, S., Malizia, F., Latora, V., Frasca, M. & Gómez-Gardeñes, J. Hyperedge overlap drives synchronizability of systems with higher-order interactions. *Phys. Rev. E* **111**, 034302 (2025).
46. Wasserman, S., Faust, K. et al. *Social network analysis: Methods and applications* (Cambridge University Press, 1994).
47. Milo, R. et al. Superfamilies of evolved and designed networks. *Science* **303**, 1538–1542 (2004).
48. Landry, N. W., Young, J.-G. & Eikmeier, N. The simpliciality of higher-order networks. *EPJ data Sci.* **13**, 17 (2024).
49. Akoglu, L., Tong, H. & Koutra, D. Graph based anomaly detection and description: a survey. *Data Min. Knowl. Discov.* **29**, 626–688 (2015).
50. Feng, S. et al. Hypergraph models of biological networks to identify genes critical to pathogenic viral response. *BMC Bioinforma.* **22**, 287 (2021).
51. Lotito, Q. F., Musciotto, F., Battiston, F. & Montresor, A. Exact and sampling methods for mining higher-order motifs in large hypergraphs. *Computing* **106**, 475–494 (2024).
52. Mancastropa, M., Iacopini, I., Petri, G. & Barrat, A. Hyper-cores promote localization and efficient seeding in higher-order processes. *Nat. Commun.* **14**, 6223 (2023).
53. Genetti, S., Ribaga, E., Cunegatti, E., Lotito, Q. F. & Iacca, G. Influence maximization in hypergraphs using multi-objective evolutionary algorithms. In *International Conference on Parallel Problem Solving from Nature*, 217–235 (Springer, 2024).
54. Kraakman, Y. J. & Stegehuis, C. Hypercurveball algorithm for sampling hypergraphs with fixed degrees. *J. Complex Netw.* **13**, cna007 (2025).
55. Chodrow, P. S. Configuration models of random hypergraphs. *J. Complex Netw.* **8**, cnaa018 (2020).
56. Barthelemy, M. Class of models for random hypergraphs. *Phys. Rev. E* **106**, 064310 (2022).
57. Nakajima, K., Shudo, K. & Masuda, N. Randomizing hypergraphs preserving degree correlation and local clustering. *IEEE Trans. Netw. Sci. Eng.* **9**, 1139–1153 (2021).
58. Saracco, F., Petri, G., Lambiotte, R. & Squartini, T. Entropy-based models to randomise real-world hypergraphs. *Commun. Phys.* **8**, 284 (2025).
59. Musciotto, F., Battiston, F. & Mantegna, R. N. Detecting informative higher-order interactions in statistically validated hypergraphs. *Commun. Phys.* **4**, 1–9 (2021).
60. Lotito, Q. F. et al. Hypergraphx: a library for higher-order network analysis. *J. Complex Netw.* **11**, cna019 (2023).

## Acknowledgements

F.B. acknowledges support from the Air Force Office of Scientific Research under award number FA8655-22-1-7025. F.B. acknowledges support from the Austrian Science Fund (FWF) through project 10.55776/PAT1052824 and project 10.55776/PAT1652425. A.M. acknowledges support from the European Union through Horizon Europe CLOUDSTARS project (101086248).

## Author contributions

Q.F.L., A.M., and F.B. designed research; Q.F.L. analyzed data; Q.F.L. and A.V. developed the algorithms and performed the computations; Q.F.L., A.V., A.M., and F.B. analyzed results and wrote the paper.

## Competing interests

Federico Battiston is an Editorial Board Member for *Communications Physics*, and one of the editors curating the Focus Collection “Higher-order Interaction Networks 2024” but was not involved in the editorial review of, or

the decision to publish this article. All other authors declare no competing interests.

### Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s42005-025-02472-9>.

**Correspondence** and requests for materials should be addressed to Quintino Francesco Lotito or Federico Battiston.

**Peer review information** : *Communications Physics* thanks the anonymous reviewers for their contribution to the peer review of this work.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026