

Identifying maximal sets of significantly interacting nodes in higher-order networks

Lorenzo Betti ^{1,*}, Federico Musciotto ^{2,*}, Federico Battiston ^{1,3} and Rosario N. Mantegna ^{2,4}

¹*Department of Network and Data Science, Central European University, 1100 Vienna, Austria*

²*Dipartimento di Fisica e Chimica Emilio Segrè, Università di Palermo, Viale delle Scienze, Ed. 18, I-90128 Palermo, Italy*

³*Department of AI, Data and Decision Sciences, Luiss University of Rome, Viale Romania 32, 00197 Rome, Italy*

⁴*Complexity Science Hub Vienna, Josefstädter Strasse 39, 1080 Vienna, Austria*



(Received 16 February 2026; accepted 30 June 2026; published 26 August 2026)

Filtering methods are fundamental tools for extracting the backbone of complex networks. Systems displaying group interactions, however, open the way to new types of filtering approaches. Here, we introduce a statistical filtering method for higher-order networks that identifies statistically validated maximal interacting sets—maximal sets of nodes that consistently interact together within group interactions. Using properly designed benchmarks, we show that our approach is highly effective in systems where the maximal sets are likely to be diluted into interactions of larger sizes that include occasional participants. Applications to real-world data reveal that the identified sets of nodes are characterized by higher levels of similarity and topical coherence, highlighting the ability of our method to provide new insights on the organization of real-world higher-order networks.

DOI: [10.1103/t86l-31d5](https://doi.org/10.1103/t86l-31d5)

I. INTRODUCTION

Over the past decades, the amount and complexity of data describing natural and social systems have increased [1]. Networks, which are collections of nodes and edges linking pairs of them, have emerged as a powerful tool to effectively study such systems, being able to describe the interaction patterns that shape many natural and human phenomena [2]. In many real-world systems, however, interactions can involve more than two nodes at a time. Higher-order networks have emerged as a natural generalization of pairwise networks to encode such many-body interactions [3–5], where hyperedges represent interactions of any size. This framework has been successfully applied to a variety of domains such as social [6], ecological [7], and biochemical systems [8,9].

Filtering methods have emerged as a fundamental tool to simplify the representation of overly complex networked systems. This task consists of pruning nodes or edges to reduce the amount of information while still maintaining the core properties of the system. This step is required for many practical applications such as reducing connections due to noise [10], focusing on the relevant part of the network responsible for its functioning [11], improving computation [12], or making visualization of large networks feasible [13]. Many methods have been proposed for this task, which go under the name of backboning or network sparsification [14–16], and have been applied to a variety of different domains [17–20].

These filtering methods mainly belong to two families. The first relies on imposing topological constraints that force the reduction of the number of connections in the filtered system. The minimum spanning tree belongs to this family, where a system of n connected nodes is reduced to a tree with $n - 1$ links [21], along with extensions involving new topological constraints [22,23] and statistical approaches [24,25]. The second family is based on dropping connections that are compatible with an appropriate null model [10,14,15]. These models generate ensembles of randomized replicas of the system while preserving some key features of the original network. An example is the disparity filter, which prunes edges based on their weights [14]. This statistical filter, along with similar methods [10,15], naturally accounts for the heterogeneity and multiscale connectivity pattern observed in real-world networks [26].

In higher-order networks, developing powerful methods for filtering represents an even more pressing need, given the richness of information encoded in multibody interactions [27,28]. The multibody nature of higher-order interactions not only requires one to adapt existing methods developed for pairwise networks but also opens the door to entirely new methodological approaches that are not direct generalizations of existing tools. A stream of methods—falling under the name of reducibility—have approached the issue of dimensionality reduction of higher-order systems by separating interactions according to their order (number of interacting units per interaction), assessing which interaction orders are more relevant at a global level, and retaining or pruning entire interaction orders based on their structural [29] or dynamical [30] relevance. Other approaches take a local perspective, more similar to traditional filtering approaches, such as the statistically validated hypergraphs (SVH), which, for each interaction size, validate statistically each single hyperedge based on a local null model [27]. These approaches act by

*These authors contributed equally to this work.

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

pruning existing hyperedges, with the resulting higher-order network composed exclusively of a subset of the interactions that were present originally. However, when interactions involve multiple nodes, the internal composition of a hyperedge becomes itself an object of interest. This consideration opens the door to the possibility of identifying significantly interacting sets of nodes *within* observed hyperedges, possibly across orders of interactions, in contrast to validating entire hyperedges only. Filtering approaches taking this new and broader perspective are currently still lacking.

In this work, we propose a filtering methodology for higher-order complex systems capable of extracting Statistically validated maximal interacting sets (SVMIS) from collectively interacting groups of nodes across interactions of different sizes. Our method relies on a null model designed for repeated higher-order interactions that is able to identify the maximal sets of nodes that consistently interact collectively and discard the nodes that co-interact with them only occasionally, while inheriting the capability of preserving network heterogeneity typical of other network filtering approaches [15,27]. Given the combinatorial nature of the problem at hand, we propose an approximation method to calculate the statistical significance of sets of nodes that overcomes this computational bottleneck (see Sec. S1 of the Supplemental Material [31]). We assess the method's performance on a class of synthetic benchmarks and demonstrate its ability to extract meaningful structure in empirical data, highlighting its general applicability to higher-order complex systems.

II. FILTERING METHOD

Given a collection of group interactions, we aim at developing a method that identifies sets of nodes that consistently interact with one another across collective interactions. In other words, we want to identify those sets of nodes that repeatedly appear together *within* group interactions more often than expected by chance under an appropriate null model. Unlike traditional filtering methods that prune nodes or entire interactions, our method focuses on filtering nodes *within* group interactions to reveal statistically meaningful sets of interacting nodes. For example, in a retail setting where group interactions correspond to shopping baskets, consider a set of m products that are functionally related, such as yeast, flour, mozzarella, and tomato sauce, which are needed to make pizza. Although these products may often be bought along with other, potentially unrelated products because of the heterogeneity of shopping habits, classic filtering methods would merely prune *entire* shopping baskets, thus overlooking the recurring co-occurrence of sets of items *within baskets*. Instead, our method aims at detecting such sets of m products that systematically co-appear in baskets of size m or larger, out of all the ones that are only occasionally co-purchased with them.

In the following we adopt the formalism of higher-order networks [3] to develop our method. Within this formalism, we represent systems as hypergraphs, where higher-order interactions between an arbitrary number of nodes are encoded as hyperedges. A hyperedge of size n corresponds to an interaction among n nodes (e.g., n products in the same basket). Our purpose can be framed as transforming a collection of




FIG. 1. Statistically validated sets. Hyperedges (shaded areas) that contain a statistically validated set of nodes (blue nodes). The three validated nodes consistently appear together in different hyperedges, which also contain other nodes.

hyperedges into a set of statistically validated sets, where each of them represents a group of consistently interacting nodes (Fig. 1).

It is worth stressing that the validated node sets are not necessarily a subset of the hyperedges of the original hypergraph because m nodes consistently appearing together may always occur within hyperedges of size m or larger. This aspect distinguishes our method from more traditional filters developed for higher-order networks like SVH [27] and other related approaches [28,32]. Additionally, our method is also related to the notion of nestedness in hypergraphs, which quantifies the extent to which hyperedges are entirely embedded within larger ones [33–36]. In nested systems, subsets of nodes that repeatedly appear across multiple interaction sizes are natural candidates to emerge as validated interacting sets of nodes. From this perspective, our method can be interpreted as identifying the local node-level cores that underpin broader nestedness patterns. Conversely, significantly interacting nodes can also arise in the absence of strict containment relations between hyperedges, as illustrated in Fig. 1, showing that our method and measures of nestedness capture complementary aspects of the organization of higher-order networks.

In what follows, we will call the filtered sets of interacting nodes statistically validated maximal interacting sets (SVMIS).

More concretely, our method works as follows. We consider a hypergraph with N hyperedges and, given a set of m nodes $S = \{i_1, i_2, \dots, i_m\}$, we define $N_{i_1 i_2 \dots i_m}$ as the number of times these nodes co-appear in hyperedges of size m or larger. In the case of a weighted hypergraph, where edge weights represent the multiplicity of group interactions, $N_{i_1 i_2 \dots i_m}$ is defined by accounting for these multiplicities. Let N_i also be the number of hyperedges in which node i participates, regardless of the size. Our null hypothesis states that each node i independently selects the hyperedges in which it participates uniformly at random, while keeping N_i fixed. Equivalently, in the bipartite representation of a hypergraph—where one layer refers to nodes and the other to hyperedges—this null hypothesis corresponds to a configuration-model randomization that

fixes the degree in the node layer (i.e., N_i) and rewires their connections uniformly at random [37]. The task of identifying statistically validated interacting sets of nodes corresponds to computing the probability $p(N_{i_1 i_2 \dots i_m})$ that a given set S of m nodes co-occurs in a sufficient number of hyperedges, as compared to the expectation according to the null model.

To compute the probability that nodes in S co-occur inside $N_{i_1 i_2 \dots i_m}$ hyperedges, we apply a hierarchical procedure. For the sake of illustration, we first derive the case where the set S is composed of three nodes ($m = 3$) before providing the general expression. The probability of observing the first two nodes co-occurring in exactly X_{12} hyperedges follows a hypergeometric distribution:

$$H(X_{12}|N, N_1, N_2) = \frac{\binom{N_1}{X_{12}} \binom{N-N_1}{N_2-X_{12}}}{\binom{N}{N_2}}. \quad (1)$$

Conditioned on X_{12} , the probability of observing the three nodes occurring in N_{123} hyperedges is again described by a hypergeometric distribution:

$$H(N_{123}|N, X_{12}, N_3) = \frac{\binom{X_{12}}{N_{123}} \binom{N-X_{12}}{N_3-N_{123}}}{\binom{N}{N_3}}. \quad (2)$$

Then, by summing over all feasible values of X_{12} , we obtain the probability of observing N_{123} hyperedges where the three considered nodes are all present at the same time:

$$\begin{aligned} p(N_{123}) &= \sum_{X_{12}} H(X_{12}|N, N_1, N_2) \times H(N_{123}|N, X_{12}, N_3) \\ &= \frac{1}{\binom{N}{N_2} \binom{N}{N_3}} \sum_{X_{12}} \binom{N_1}{X_{12}} \binom{N-N_1}{N_2-X_{12}} \binom{X_{12}}{N_{123}} \binom{N-X_{12}}{N_3-N_{123}}, \end{aligned} \quad (3)$$

where $X_{12} \in [N_{123}, \min(N_1, N_2)]$. To obtain the probability for an arbitrary number of nodes m , we iterate further the procedure and we get

$$\begin{aligned} p(N_{12\dots m}) &= \sum_{X_{12}} H(X_{12}|N, N_1, N_2) \\ &\quad \times \sum_{X_{123}} H(X_{123}|N, X_{12}, N_3) \times \dots \\ &\quad \dots \times \sum_{X_{1\dots m-1}} H(X_{1\dots m-1}|N, X_{1\dots m-2}, N_{m-1}) \\ &\quad \times H(N_{1\dots m}|N, X_{1\dots m-1}, N_m), \end{aligned} \quad (4)$$

where $X_{12\dots r}$ denotes the number of hyperedges containing the first r nodes and the summation limits are given by

$$X_{12\dots r} \in [N_{1\dots r+1}, \min(X_{1\dots r-1}, N_r)]. \quad (5)$$

Note that Eq. (4) is symmetrical for permutations of the order of the nodes, making this expression suitable for any set of nodes regardless of the order and thus $p(N_{12\dots m}) = p(N_{i_1 i_2 \dots i_m})$ for any permutation of the nodes in S . Although Eq. (4) may appear similar to Eq. (3) in the reference introducing the SVH method [27], there is a crucial difference. For SVH, since hyperedges are validated for each order separately, the null expectation is built by considering all hyperedges of size m . Instead, in our approach we look at all hyperedges of size

m or larger, thus making it possible to filter nodes *within* hyperedges. In summary, Eq. (4) is an iterative composition of hypergeometric intersections [38].

Having derived an expression for $p(N_{i_1 i_2 \dots i_m})$, we can finally assess the statistical significance of a set of m nodes against a threshold of statistical significance α . The corresponding p value of the given set of nodes can be obtained through the survival function:

$$p(X \geq N_{i_1 i_2 \dots i_m}) = 1 - \sum_{x=0}^{N_{i_1 i_2 \dots i_m}-1} p(x). \quad (6)$$

For the rest of this paper, we set $\alpha = 0.01$. We apply a Benjamini-Hochberg correction for false discovery rate to account for the large number of tested hypotheses [39] (see Sec. S2 of the Supplemental Material [31] for details).

Since we are interested in finding statistically validated maximal interacting sets, we are in principle interested in assessing all possible combinations of nodes of any size. This approach would be computationally unfeasible because the number of such sets grows as the Bell number of the number of nodes (i.e., faster than exponential). However, since sets of nodes that never appear together in any hyperedge cannot represent meaningful sets of interacting nodes [$N_{i_1 i_2 \dots i_m} = 0$ making $p(X \geq N_{i_1 i_2 \dots i_m}) = 1$], we restrict the set of nodes to test to those that appear at least once together within a hyperedge (i.e., such that $N_{i_1 i_2 \dots i_m} \geq 1$), which reduces the search space. Moreover, we focus on maximal interacting sets. Namely, once a set of m nodes is found to be statistically validated, we do not test any of the subsets of such m nodes. Note that those subsets of a statistically validated set should not necessarily be considered as significantly interacting as well. Thus our approach mainly focuses on the maximally large set of interacting nodes and can thus neglect other smaller significant sets within larger ones.

To further ease this computational bottleneck, we propose an approximation for the computation of the p value by approximating Eq. (4) with a binomial distribution, which tends to overestimate the exact p value (see Sec. S1 of the Supplemental Material [31] for details). This approximation recalls the one proposed in Ref. [40]. In all the applications that we are going to showcase in the following, we will use the approximated p value.

Given those considerations, for each size m , we proceed by testing sets of nodes from largest to smallest, where the largest m corresponds to the size of the largest observed hyperedge.

III. STATISTICALLY VALIDATED MAXIMAL INTERACTING SETS

A. Definition and investigation of benchmarks

To test the effectiveness of our approach in identifying sets of nodes that interact frequently within hyperedges, we design benchmark hypergraphs with implanted sets of nodes representing the interacting groups that the method should recover. Starting from these implanted sets, we generate hyperedges by eventually “diluting” them with additional nodes selected at random. This controlled setting is not intended to reproduce all structural features of empirical hypergraphs, but rather to showcase the context in which SVMIS is expected to

be useful: identifying recurrent core groups that may appear within larger higher-order interactions.

We create benchmark hypergraphs made of N nodes in which we implant T sets for each size n . The number of implanted sets is controlled through a density parameter d such that, for each realization of the benchmark, we set $T = d \binom{N}{n}$ to make the number of implanted sets grow with N while keeping benchmarks of different sizes comparable. In this work, we explored $N = \{100, 200, 500\}$ and $d \in \{0.005, 0.01, 0.02\}$, with n ranging in the interval $[2, 4]$. For each implanted set, we create l hyperedges of size up to $n_{\max}$, where l is sampled from a binomial distribution with $p_{\text{bin}} = 0.5$ and $n_{\text{bin}} = 20$. Each hyperedge contains the n nodes of the implanted set plus n' additional nodes, with n' sampled from the discrete uniform distribution $U(0, n_{\max} - n)$. This sampling yields hyperedges that either coincide with the implanted sets (when $n' = 0$) or are diluted with randomly chosen nodes (when $n' > 0$). In our simulations, we set $n_{\max} = 6$ to obtain a moderate dilution of the implanted sets into larger hyperedges, thus having the maximum hyperedge size in $[8, 10]$ depending on the value of n . In this way the n nodes of a set interact l times in different hyperedges, which always contain the n core nodes plus a set of randomly selected ones. Moreover, we introduce a parameter c that controls the level of closure at each group size. Specifically, for an implanted set of size n , c represents the closure probability, defined as the probability that all lower-order interactions among the n nodes are fully nested within the group [29,33,35]. For example, when $n = 3$, the closure probability c determines the probability of including all the three possible pairwise interactions among the nodes.

We compare the results of SVMIS against the statistically validated hypergraphs (SVH) method [27]. SVH was developed for hypergraphs but it only prunes entire hyperedges, in contrast to SVMIS.

To validate our results, we compute the true positive rate (TPR), which quantifies the ratio of implanted sets that we are able to identify, and the false discovery rate (FDR), which quantifies the ratio of erroneously detected implanted sets that the different methods select. Specifically, $TPR = \frac{TP}{TP+FN}$ and $FDR = \frac{FP}{FP+TP}$, where TP are the true positives (i.e., significant sets corresponding to the originally implanted sets), FN are the false negatives (i.e., implanted sets that are not identified), and FP are the false positives (sets detected as significant by the method that are not among the originally implanted sets). We provide sensitivity analyses for different values of α in Sec. S3 of the Supplemental Material [31].

Figure 2(a) shows TPR as a function of c for $N = 200$ and $d = 0.005$, but the observed behavior is independent of the specific values of N and d . Each point is the median of 100 benchmark realizations. The color bands represent the interval between the 10th and the 90th percentile. The performance of SVMIS and SVH is not affected by c , as both are methods designed to take into account the higher-order nature of interactions. In fact, if three nodes i, j, k constitute a clique of (significant) pairwise interactions but never interact altogether, SVMIS and SVH will not statistically validate the corresponding set because $N_{ijk} = 0$ even if N_{ij} , N_{ik} , and N_{jk} are not. As a consequence, SVMIS and SVH will not miss $\{ij\}$, $\{jk\}$, and $\{ik\}$ as a significant set. We also observe that

SVMIS significantly outperforms SVH. This happens because in our benchmark hypergraphs we allow the n nodes of an implanted set to eventually interact within larger hyperedges. While SVMIS counts all the hyperedges in which the n nodes appear, the SVH approach misses all those in which there are additional, randomly selected nodes. For this reason, SVH is more likely to underestimate the significance of a set and thus have a lower value of TPR. Figure 2(b) plots the FDR as a function of c . SVH and SVMIS almost never detect wrong sets as they distinguish between interactions of different orders.

To better characterize the difference in performance between SVMIS and SVH, we modify the benchmark by enforcing the control on the extent to which the implanted sets are diluted within larger hyperedges. We introduce a parameter f that regulates the fraction of times in which the n nodes of an implanted set appear together with additional nodes (i.e., in hyperedges of size larger than n). Specifically, for each implanted set of size n , we create l hyperedges such that (i) a fraction f of newly created hyperedges are generated by diluting the implanted set with other n' random nodes, with n' sampled from $U(1, n_{\max} - n)$, and (ii) the remaining fraction $1 - f$ consists of hyperedges equal to the implanted set itself (i.e., $n' = 0$). The change on the lower bound of n' from 0 to 1 ensures that f fully controls the fraction of implanted sets diluted within larger hyperedges. The larger f , the more we expect SVH to underestimate the number $N_{i_1 i_2 \dots i_m}$ of collective interactions and thus will likely fail to validate the true sets. In Fig. 2(c) we show TPR as a function of f . Indeed, we see that SVH significantly worsens its performance when f increases. Figure 2(d) shows the behavior of FDR as a function of f . The FDR for SVH remains low not because SVH is accurate, but because the method detects almost no positives [see the TPR in Fig. 2(c)].

It is worth noting that the TPR values shown in Fig. 2 represent a lower bound on the performance achievable with the exact p -value computation. Since the approximate p values tend to overestimate the exact ones (see Sec. S1 of the Supplemental Material [31]), the approximation exclusively introduces additional false negatives (i.e., implanted sets that fail to be validated) without affecting the false discovery rate. The consistently high TPR observed across all benchmark settings therefore confirms that the approximate p -value computation offers a favorable trade-off: it makes SVMIS computationally feasible while preserving strong performance in identifying significantly interacting sets of nodes within higher-order interactions.

In sum, the benchmark analysis reveals that SVMIS is able to identify sets of significantly interacting nodes within larger group interactions. Moreover, the comparison against SVH confirms that these filtering approaches differ from SVMIS as a result of the different problem they aim to address. In the following, we apply SVMIS to real-world data.

B. Application to real data: Baskets of co-purchased goods at Walmart

We analyze a data set that tracks 69 906 co-purchases of 88 860 products at Walmart [41], available as part of the data repository Hypergraphx-data [42]. This system can be mapped into a hypergraph where hyperedges represent




FIG. 2. Benchmark analysis. (a) True positive rate (TPR, y axis) between the implanted sets and those detected by statistically validated hypergraphs (SVH, sea-green line) and statistically validated maximal interacting sets (SVMIS, orange line) as a function of the closure probability c with $N = 200$ and $d = 0.005$. Solid lines refer to the median across 100 realizations of the benchmark and transparent bands to the interval between the 10th and the 90th percentile. (b) False discovery rate (FDR, y axis) for SVH and SVMIS as a function of c for the same benchmark realizations. (c) TPR (y axis) for SVH and SVMIS as a function of the fraction of larger hyperedges f over 100 benchmark realizations. (d) FDR (y axis) for SVH and SVMIS as a function of the fraction of larger hyperedges f over the same benchmark realizations of (c).

co-purchases. In this context, we use the SVMIS approach to extract sets of products that are inherently bought together. When extracting the statistically validated maximal interacting sets we exclude from the hypergraph all baskets of size larger than 10 or equal to 1. Indeed, each large basket of size k which is not statistically significant produces a large number of candidate sets of smaller size k' —specifically it produces $\binom{k}{k'}$ candidates. Given that larger baskets have a small number of co-occurrence, we cut the system at size 10 in order to avoid the explosion in the number of hyperedges with low chance of getting validated.

Figure 3(a) shows the number of validated (blue) and not validated (orange) sets for each size. The number of validated sets exponentially decays with size and, for large n , it is several orders of magnitude smaller than the sets that are observed at least once but are not rejecting the null hypothesis.

The data set also reports the category of each product. The categories are as follows: clothing, shoes, and acces-

sories; electronics and office; home, furniture, and appliances; home improvement and patio; baby; toys, games, and video games; food, household, and pets; pharmacy, health, and beauty; sports, fitness, and outdoors; auto, tires, and industrial; other. By using this information, we define a proxy of “purity” of co-purchased goods. Purity is a measure that goes from 0 (if all products in a group belong to different categories) to 1 (if they all belong to the same category) and is defined as $P = \frac{n}{n-1}(1 - \frac{C}{n})$, where n is the size of the set and C is the number of unique categories to which the n products belong. Figure 3(b) shows the “purity” for the validated (blue) and not validated (orange) sets as a function of their size. In the figure, the solid lines plot the mean and the transparent bands the error of the mean. These results show that, for each size, the validated sets identified by SVMIS correspond to sets of items that are more homogeneous with respect to category than the ones that are not validated.




FIG. 3. Significant co-purchases at Walmart. (a) Counts of validated sets (blue bars) and not validated ones (orange bars) as a function of size. The sum of blue and orange bars constitutes the total number of tested sets. (b) Solid lines represent the average purity and transparent bands are the standard error for validated (blue) and not validated (orange) sets.

C. Co-traded stocks belonging to the S&P100 index

We investigate a data set tracking tick-by-tick transactions performed by 96 stocks that belong to the S&P100 index and are traded on any US venue in the month 12/2020 [43]. The data comes with a temporal resolution of 1 ms. By applying our method to this data set, we detect SVMISs of stocks that are co-traded in the investigated period in excess to what is expected by a null model taking into account the heterogeneity observed in stock trading. Again, before extracting the SVMISs we exclude from the hypergraph all hyperedges larger than 10 stocks or equal to 1. We perform our analysis at a daily scale and then aggregate the results obtained for different days.

Figure 4(a) shows the number of validated (blue) and not validated (orange) sets for each size. Figure 4(b) shows the mean correlation (solid line) and its error (transparent bands) between couples of stocks in validated (blue) and not validated (orange) sets. Correlation is computed on time series of financial returns extracted at a resolution of 1 min.

By identifying sets of stocks that are co-traded significantly more often than expected by chance, our method reveals higher-order structures in price formation characterized by

significantly higher average price correlations compared to nonvalidated sets. This provides empirical confirmation to the theory that high-frequency collective activity contributes significantly to price synchronization across related assets [44], linking the algorithmic driven co-occurrence of market interactions directly to the emergence of coordinated price dynamics.

D. Food recipes

Finally, we apply our method on a data set representing food recipes from different cuisines around the world [45]. We used the preprocessed data set provided in Ref. [46], which can be represented as a hypergraph where ingredients correspond to nodes and recipes (i.e., groups of ingredients) to hyperedges. After excluding recipes composed of one single ingredient or more than 10 ingredients, the resulting hypergraph has 580 nodes and 32 291 hyperedges.

Figure 5(a) shows the number of validated (blue) and not validated (orange) sets as a function of their size. Validated sets are concentrated at intermediate sizes and display a rapid decay for larger sizes. This suggests that world cuisines are structured around relatively small and recurrent combinations of ingredients.




FIG. 4. Co-traded stocks from S&P100 index. (a) Counts of validated sets (blue bars) and not validated ones (orange bars) as a function of size. The sum of blue and orange bars constitutes the total number of tested sets. (b) Solid lines represent the average correlation and transparent bands are the standard error among stocks in the validated (blue) and not validated (orange) sets.




FIG. 5. Food recipes. (a) Counts of validated sets (blue bars) and not validated ones (orange bars) as a function of size. The sum of blue and orange bars constitutes the total number of tested sets. (b) Frequency-rank plots of ingredients across recipes (violet line), validated sets across real recipes (red line), and across randomized recipes (gray lines). The randomized ones represent the result from 50 independent realizations of the randomized model.

To illustrate the nontrivial nature of the validated sets identified by SVMIS, we examine their frequency across recipes. Figure 5(b) shows the frequency-rank plot of validated sets (red line), which spans almost three orders of magnitude and exhibits finite-range scaling behavior. In contrast, the frequency-rank plot of ingredients is distributed more narrowly (violet line).

We compare this observation against a simple randomized model where, for each cuisine, we assign ingredients to recipes randomly while preserving ingredient usage frequency and the distribution of the number of ingredients per recipe (i.e., a bipartite configuration model in the ingredient-recipe bipartite network). Although displaying a finite-range scaling behavior, the slope is steeper compared to the one observed in the real system (gray lines).

In sum, these observations illustrate that SVMIS identifies nontrivial recurrent structures in the recipe hypergraph, which are not captured by mere ingredient frequency statistics and are not compatible with random combination of ingredients. In this sense, the method reveals nontrivial organization at the level of ingredient combinations.

III. DISCUSSION

The recent surge in interest towards higher-order networks, together with the increased availability of a large amount of data on several real world systems, have led to the introduction of the first filtering tools for hypergraphs [27,40]. However, the task of detecting statistically validated units of collectively interacting elements is multifaceted and the multidimensional nature of higher-order interactions calls for the design of different validation approaches. The existing methods have focused on the extension of validating approaches from pair-

wise interactions to hyperedges [40] and the detection of statistically validated group interactions that occur only at a specific size [27]. Here, we introduced a different and more general framework for the detection of statistically validated maximal interacting sets, which represents sets of nodes that appear together within group interactions. The computational efforts needed for the statistical validation are demanding but accessible with ordinary PCs for groups up to a size of about 10 elements. To speed up the computation we have proposed an approximated procedure that works well in a large number of cases as shown in our investigation of benchmarks. Using properly designed hypergraph benchmarks, we show that our approach is highly effective on higher-order systems (i) in which informative interactions among n core nodes are likely to be diluted into larger hyperedges that feature occasional nodes and (ii) with a high degree of closure. We also show how, on real data, our approach allows us to detect significantly interacting sets whose nodes are characterized by high levels of similarity, which might help in shedding light on the generative process of the underlying hypergraphs.

ACKNOWLEDGMENTS

F.B. acknowledges support from the Austrian Science Fund (FWF) through Projects No. 10.55776/PAT1052824 and No. 10.55776/PAT1652425.

DATA AVAILABILITY

Data and scripts needed to replicate the results are available at the GitHub repository [47]. The code to extract statistically validated maximal interacting sets is available as part of the library Hypergraphx [48].

- [1] M. E. J. Newman, The structure and function of complex networks, *SIAM Rev.* **45**, 167 (2003).
- [2] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D.-U. Hwang, Complex networks: Structure and dynamics, *Phys. Rep.* **424**, 175 (2006).
- [3] F. Battiston, G. Cencetti, I. Iacopini, V. Latora, M. Lucas, A. Patania, J.-G. Young, and G. Petri, Networks beyond pairwise interactions: Structure and dynamics, *Phys. Rep.* **874**, 1 (2020).
- [4] F. Battiston, E. Amico, B. Alain, *et al.*, The physics of higher-order interactions in complex systems, *Nat. Phys.* **17**, 1093 (2021).
- [5] F. Battiston, V. Capraro, F. Karimi, S. Lehmann, A. B. Migliano, O. Sadekar, A. Sánchez, and M. Perc, Higher-order interactions shape collective human behaviour, *Nat. Hum. Behav.* **9**, 2441 (2025).
- [6] A. Patania, G. Petri, and F. Vaccarino, The shape of collaborations, *EPJ Data Sci.* **6**, 18 (2017).
- [7] J. Grilli, G. Barabás, M. J. Michalska-Smith, and S. Allesina, Higher-order interactions stabilize dynamics in competitive network models, *Nature (London)* **548**, 210 (2017).
- [8] J. Jost and R. Mulas, Hypergraph Laplace operators for chemical reaction networks, *Adv. Math.* **351**, 870 (2019).
- [9] G. Petri, P. Expert, F. Turkheimer, R. Carhart-Harris, D. Nutt, P. J. Hellyer, and F. Vaccarino, Homological scaffolds of brain functional networks, *J. R. Soc. Interface* **11**, 20140873 (2014).
- [10] M. Coscia and F. M. H. Neffke, Network backbone with noisy data, in *2017 IEEE 33rd International Conference on Data Engineering (ICDE)* (IEEE, New York, 2017), pp. 425–436.
- [11] D. Grady, C. Thiemann, and D. Brockmann, Robust classification of salient links in complex networks, *Nat. Commun.* **3**, 864 (2012).
- [12] A. Mercier, S. Scarpino, and C. Moore, Effective resistance against pandemics: Mobility network sparsification for high-fidelity epidemic simulations, *PLoS Comput. Biol.* **18**, e1010650 (2022).
- [13] M. Imre, J. Tao, Y. Wang, Z. Zhao, Z. Feng, and C. Wang, Spectrum-preserving sparsification for visualization of big graphs, *Comput. Graphics* **87**, 89 (2020).
- [14] M. Á. Serrano, M. Boguná, and A. Vespignani, Extracting the multiscale backbone of complex weighted networks, *Proc. Natl. Acad. Sci. USA* **106**, 6483 (2009).
- [15] M. Tumminello, S. Micciche, F. Lillo, J. Piilo, and R. N. Mantegna, Statistically validated networks in bipartite complex systems, *PLoS One* **6**, e17994 (2011).
- [16] A. Kirkley, Fast nonparametric inference of network backbones for weighted graph sparsification, *Phys. Rev. X* **15**, 031013 (2025).
- [17] M. Tumminello, F. Lillo, J. Piilo, and R. N. Mantegna, Identification of clusters of investors from their real trading activity in a financial market, *New J. Phys.* **14**, 013041 (2012).
- [18] F. Musciotto, L. Marotta, S. Micciché, J. Piilo, and R. N. Mantegna, Patterns of trading profiles at the Nordic stock exchange. A correlation-based approach, *Chaos Solitons Fractals* **88**, 267 (2016).
- [19] F. Musciotto, L. Marotta, J. Piilo, and R. N. Mantegna, Long-term ecology of investors in a financial market, *Palgrave Commun.* **4**, 92 (2018).
- [20] M.-X. Li, V. Palchykov, Z.-Q. Jiang, K. Kaski, J. Kertész, S. Micciche, M. Tumminello, W.-X. Zhou, and R. N. Mantegna, Statistically validated mobile communication networks: The evolution of motifs in European and Chinese data, *New J. Phys.* **16**, 083038 (2014).
- [21] R. N. Mantegna, Hierarchical structure in financial markets, *Eur. Phys. J. B* **11**, 193 (1999).
- [22] M. Tumminello, T. Aste, T. Di Matteo, and R. N. Mantegna, A tool for filtering information in complex systems, *Proc. Natl. Acad. Sci. USA* **102**, 10421 (2005).
- [23] G. P. Massara, T. Di Matteo, and T. Aste, Network filtering for big data: Triangulated maximally filtered graph, *J. Complex Networks* **5**, 161 (2016).
- [24] M. Tumminello, C. Coronnello, F. Lillo, S. Micciche, and R. N. Mantegna, Spanning trees and bootstrap reliability estimation in correlation-based networks, *Int. J. Bifurcation Chaos* **17**, 2319 (2007).
- [25] F. Musciotto, L. Marotta, S. Micciché, and R. N. Mantegna, Bootstrap validation of links of a minimum spanning tree, *Physica A* **512**, 1032 (2018).
- [26] A. Barrat, M. Barthélemy, R. Pastor-Satorras, and A. Vespignani, The architecture of complex weighted networks, *Proc. Natl. Acad. Sci. USA* **101**, 3747 (2004).
- [27] F. Musciotto, F. Battiston, and R. N. Mantegna, Detecting informative higher-order interactions in statistically validated hypergraphs, *Commun. Phys.* **4**, 218 (2021).
- [28] N. W. Landry, I. Amburg, M. Shi, and S. G. Aksoy, Filtering higher-order datasets, *J. Phys.: Complexity* **5**, 015006 (2024).
- [29] A. Kirkley, H. Felipe, and F. Battiston, Structural reducibility of hypergraphs, *Phys. Rev. Lett.* **135**, 247401 (2025).
- [30] M. Lucas, L. Gallo, A. Ghavasieh, F. Battiston, and M. De Domenico, Reducibility of higher-order networks from dynamics, *Nat. Commun.* **17**, 1551 (2026).
- [31] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/t861-31d5> for details on the approximate p -value computation, the Benjamini-Hochberg correction, and the sensitivity analysis to α .
- [32] A. Ceria and F. W. Takes, The relevance of higher-order ties, *EPJ Data Sci.* **14**, 62 (2025).
- [33] Q. F. Lotito, F. Musciotto, A. Montresor, and F. Battiston, Higher-order motif analysis in hypergraphs, *Commun. Phys.* **5**, 79 (2022).
- [34] T. LaRock and R. Lambiotte, Encapsulation structure and dynamics in hypergraphs, *J. Phys.: Complexity* **4**, 045007 (2023).
- [35] N. W. Landry, J.-G. Young, and N. Eikmeier, The simpliciality of higher-order networks, *EPJ Data Sci.* **13**, 17 (2024).
- [36] S. Lamata-Otín, F. Malizia, V. Latora, M. Frasca, and J. Gómez-Gardeñes, Hyperedge overlap drives synchronizability of systems with higher-order interactions, *Phys. Rev. E* **111**, 034302 (2025).
- [37] T. Squartini, J. de Mol, F. den Hollander, and D. Garlaschelli, Breaking of ensemble equivalence in networks, *Phys. Rev. Lett.* **115**, 268701 (2015).
- [38] M. Wang, Y. Zhao, and B. Zhang, Efficient test and visualization of multi-set intersections, *Sci. Rep.* **5**, 16923 (2015).
- [39] Y. Benjamini and Y. Hochberg, Controlling the false discovery rate: A practical and powerful approach to multiple testing, *J. R. Stat. Soc. B* **57**, 289 (1995).
- [40] T. Kobayashi, T. Takaguchi, and A. Barrat, The structured backbone of temporal social ties, *Nat. Commun.* **10**, 220 (2019).

- [41] I. Amburg, N. Veldt, and A. Benson, Clustering in graphs and hypergraphs with categorical edge labels, in *Proceedings of The Web Conference 2020, WWW '20* (ACM, New York, 2020), pp. 706–717.
- [42] Q. F. Lotito, L. Betti, B. Nortier, A. Montresor, and F. Battiston, Hypergraphx-data: A repository for higher-order network data, *J. Complex Networks* **14**, cnag014 (2026).
- [43] FirstRate Data, Firstrate data — Historical intraday market price data, <https://firstratedata.com/>.
- [44] A. Gerig, High-frequency trading synchronizes prices in financial markets, 2015, <http://dx.doi.org/10.2139/ssrn.2173247>.
- [45] CulinaryDB, <https://cosylab.iiitd.edu.in/culinarydb/>, accessed 2026.
- [46] C. Caprioli, S. Kulkarni, F. Battiston, I. Iacopini, A. Santoro, and V. Latora, The networks of ingredient combinations as culinary fingerprints of world cuisines, *npj Sci. Food* **9**, 242 (2025).
- [47] <https://github.com/musci8/SVH>.
- [48] Q. F. Lotito, M. Contisciani, C. De Bacco, L. Di Gaetano, L. Gallo, A. Montresor, F. Musciotto, N. Ruggeri, and F. Battiston, Hypergraphx: A library for higher-order network analysis, *J. Complex Networks* **11**, cnad019 (2023).