

<https://doi.org/10.1038/s41538-025-00588-4>

The networks of ingredient combinations as culinary fingerprints of world cuisines



Claudio Caprioli^{1,11}, Saumitra Kulkarni^{2,3,11}, Federico Battiston⁴, Iacopo Iacopini^{5,6,11} ✉, Andrea Santoro^{7,8,11} ✉ & Vito Latora^{1,9,10,11}

Investigating how different ingredients are combined in popular dishes is crucial to uncover the principles behind food preferences. Here, we use data from public food repositories and network analysis to characterize and compare worldwide cuisines. Ingredients are first grouped into broader types, and each cuisine is then represented as a network in which nodes correspond to ingredient types and weighted links describe how frequently pairs of types co-occur in recipes. Cuisines differ not only in the popularity of ingredient types and range of recipe sizes, but also in the structural organization of ingredient-type combinations. By analyzing these networks, we uncover distinctive patterns of type associations that serve as culinary fingerprints. For example, European cuisines typically distribute ingredients across different types, whereas certain Asian and South American traditions emphasize one dominant type, such as vegetables or spices. The essence of these patterns is well captured by the networks' maximum spanning trees, which offer a simplified yet representative backbone for each cuisine. We demonstrate that both these full and simplified network representations enable machine learning models to identify cuisines from subsets of recipes with very high accuracy. Networks of ingredient combinations also cluster global cuisines into meaningful geo-cultural groups, reflecting shared patterns in culinary traditions. More broadly, our study offers novel insights into the structure of world cuisines, enabling data-driven approaches to their characterization, cross-cultural comparison, and potential adaptation.

The act of transforming raw ingredients into palatable sustenance, known as cooking, predates the appearance of *Homo sapiens*^{1–3}. From the earliest civilizations, recipes have served as templates, guiding individuals in the preparation of meals using specific techniques and ingredients⁴. Over millennia, culinary traditions have undergone constant refinement and expansion, reflecting the dynamic interplay of locally available ingredients with cultural influences, trade, technological advancements, and changing palates. As a result of this long and intertwined process, specific recipes have become the very backbone of diverse culinary traditions across the globe^{5,6}. While traditional recipe books have long served as repositories of culinary knowledge, the advent of computational approaches has opened new avenues for understanding recipes as algorithms⁷. Researchers have recently started leveraging vast data sets of recipes to reveal hidden patterns and

extract valuable insights into culinary practices^{8–11}. Tools coming from natural language processing¹² and machine learning¹³ have revolutionized our ability to analyze and interpret the complexities of recipe composition and preparation methods¹⁴. Applications range from extracting the list of ingredients from meal images^{15,16}, to examining nutrient concentration in food¹⁷, predicting food processing levels¹⁸ and investigating the impact of certain diets on health^{19,20}.

The essence of recipes cannot be separated from the ingredients, which as in a symphony composed of individual notes, come up together to create an unexpected harmonic consonance, each contributing its unique flavor and texture to the final dish. It is precisely this combination of ingredients — and the different techniques used to cook them — that would ultimately determine whether the resulting product would be successful or not. At the

¹Department of Physics and Astronomy, University of Catania, Catania, Italy. ²Network Science Institute, Northeastern University, Boston, MA, USA. ³Department of Scientific Computing, Modeling & Simulation, Savitribai Phule Pune University, Pune, India. ⁴Department of Network and Data Science, Central European University, Vienna, Austria. ⁵Network Science Institute, Northeastern University London, London, UK. ⁶Department of Physics, Northeastern University, Boston, MA, USA. ⁷CENTA Institute, Turin, Italy. ⁸Neuro-X Institute, École Polytechnique Fédérale de Lausanne, Geneva, Switzerland. ⁹School of Mathematical Sciences, Queen Mary University of London, London, UK. ¹⁰Complexity Science Hub, Vienna, Austria. ¹¹These authors contributed equally: Claudio Caprioli, Saumitra Kulkarni, Iacopo Iacopini, Andrea Santoro, Vito Latora. ✉e-mail: iacopo.iacopini@nulondon.ac.uk; andrea.santoro@centai.eu

heart of culinary creativity lies a quest to explore and discover novel ingredient combinations – a journey in a virtually unbounded space of pairings, mixtures, and innovations that has only been partially charted. However, not all ingredient combinations are successful, as evidenced by the considerable gap between the potential number of recipes and those recorded throughout human history²¹. This discrepancy suggests that culinary creativity often builds on top of known flavors, balancing the search for novelty with the use of established, successful combinations^{22–24}.

Complex networks^{25,26} are particularly suited to describe the way in which different ingredients are combined in recipes, and also to capture recurring patterns^{27,28}. So far, networks have been applied to describe ingredient combinations with a particular focus on their chemical composition. In particular, bipartite networks^{29–31} have been instrumental in studying the *food pairing hypothesis*, a famous culinary conjecture introduced by Chef Heston Blumenthal³², stating that the compatibility of ingredients in recipes strongly depends upon the amount of shared flavor compounds^{33–36}. Recent findings however suggest that food pairing is not a universal principle, but cuisines worldwide differ in the role chemical affinity plays in recipe formation³⁷. The growing interest in the chemical composition of food has spurred broader explorations into the implications of food consumption on health. Network-based studies have significantly advanced biomedical knowledge by unraveling the intricate interplay of nutritional value, chemical profiles, and their impact on diseases^{17,19,38–40}. It is now well understood that exploring the chemical complexity of food is an essential requirement for the study of human health, inextricably linked to the diet of individuals⁴¹. However, while analyses focused on chemical composition are pivotal in clarifying the relationship between diet and health, and at the same developing targeted therapies and precision nutrition, they often fail to capture the taste profiles of ingredients. Assortative principles based on food compounds lack generality and apply only to specific geographic cuisines. Moreover, these analyses typically overlook how culinary traditions combine ingredients in distinct and culturally meaningful ways. While reductionism is often a driving force in scientific inquiry, its application to gastronomy hides the risk of discarding valuable information encoded in culinary practices, knowledge long recognized and applied in kitchens around the world. A crucial observation is that ingredients often serve similar functional roles within recipes, regardless of their chemical makeup, such as providing umami, acidity, or aromatic complexity. For instance, soy sauce in East Asian cuisine and fish sauce in Southeast Asian cuisine both function as salty umami enhancers. Similarly, very different types of fat sources, like oil, butter, or lard, are often used as

non-stick agents and heat conductors when cooking other ingredients in a pan. This suggests that studying ingredients based on their functional associations, rather than in isolation, could reveal hidden patterns in culinary practices.


In this work, we address these limitations by categorizing ingredients into broader “food types” or “ingredient categories”. Drawing upon a recipe dataset consistently covering geographic areas from all the five continents³⁶, we employ a network-based approach to characterize and compare cuisines from all over the world. Each regional cuisine is turned into a network of ingredient combinations, whose nodes represent food types, while links describe the ways food types are combined in recipes. The analysis of such networks allows to identify unique food profiles and food co-occurrence patterns associated with each cuisine, highlighting with unprecedented precision commonalities and differences in gastronomic preferences across geographic regions.

Results

Statistical descriptors of world cuisines

We study a data set of recipes (dishes) covering 23 cuisines around the world, for a total of $R = 45661$ recipes and $I = 604$ unique ingredients (see Methods for a detailed description of the data set). This number results from a preprocessing stage aimed at standardizing the ingredient set and reducing dataset complexity, which consolidated the original 695 raw ingredient entries into 604 canonical forms. That is, ingredients with similar culinary roles or origins were grouped together, for instance, merging varieties (e.g., all pasta types as pasta), preparation forms (e.g., dried and fresh peas as peas), and subtypes (e.g., cherry and plum tomatoes as tomatoes). Additionally, ingredients produced through similar transformation processes were mapped to unified labels (e.g., all types of cheese, bread, and wine), and animal parts were aggregated by species (e.g., all chicken parts as chicken). Full details on the preprocessing are provided in the Methods and Supplementary Note 1. Each cuisine comes with its own set of recipes, where each recipe is a different combination of ingredients. Starting from the individual ingredients, we investigate how the number of ingredients and the number of recipes vary across the different cuisines. In Fig. 1, we report the total number of recipes as a function of the total number of ingredients in a lin-log scale. The distribution of distinct ingredients varies significantly among cuisines, ranging from the 144 of the Portuguese cuisine to the 503 of the American one (Acknowledging the ambiguous use of this term, we adopt the standard terminology where American refers to dishes prepared in the United

Fig. 1 | Scaling of recipes vs total ingredients across world cuisines. Relationship between the number of recipes and the total number of individual ingredients across various cuisines. Each data point represents a cuisine and is color-coded according to the continent affiliation, as illustrated in the world map reported in the top-left corner. A lin-log scale has been used in the plot, and the reported best-fit line highlights the exponential growth of the number of recipes with the total number of ingredients.



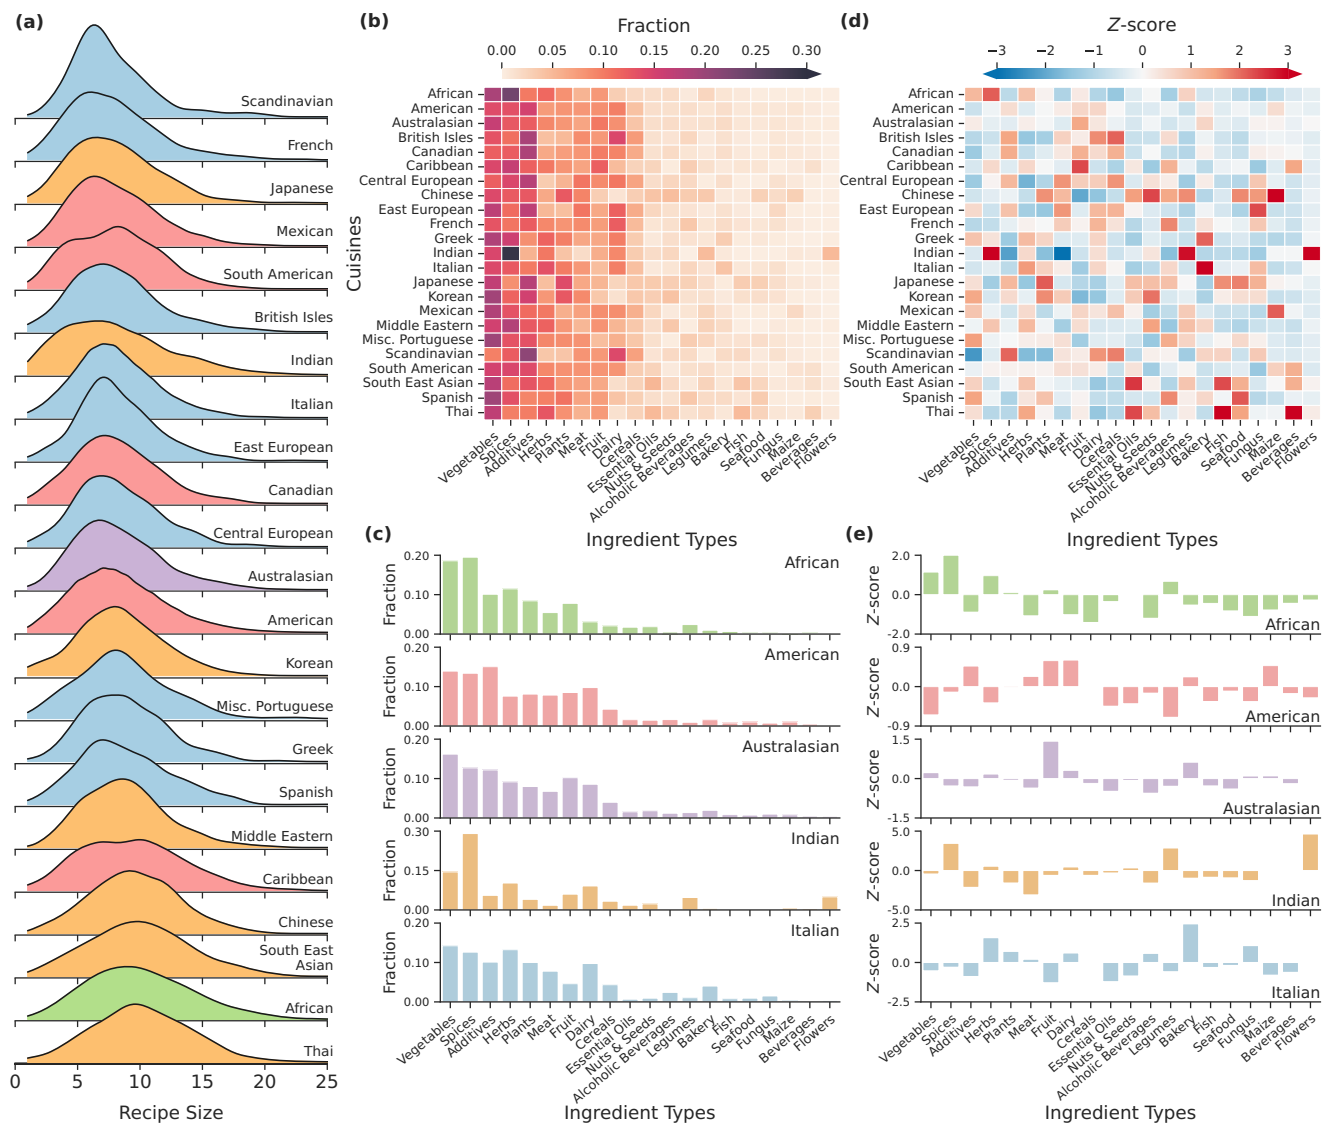


Fig. 2 | Recipe sizes and ingredient popularity. **a** Ridge plot illustrating the distribution of recipe sizes across cuisines from different continents. Each distribution, sorted by average, corresponds to a cuisine, and is colored according to the continent affiliation (as per Fig. 1). **b** Heatmap depicting the ingredient type popularity profiles of cuisines. Each row corresponds to a cuisine, and each column to an ingredient type. The color intensity represents the popularity of an ingredient type in a given cuisine. **c** Bar chart representation of ingredient type popularity profiles derived

from (b) for five specific cuisines, highlighting the variations in ingredient preferences. **d** Heatmap illustrating the z-scores of the ingredient type popularity profiles, providing insights into their relative significance across cuisines. Dark-red and dark-blue regions indicate respectively higher and lower z-scores. **e** Bar chart representation of the z-scores derived from (d) for the same five cuisines, revealing the relative significance of ingredient types.

States only (https://en.wikipedia.org/wiki/American_cuisine). Notably, the combination of ingredients in recipes differs widely among cuisines, with cuisine sizes spanning three orders of magnitude. The smallest and largest recipe data sets are again represented by Portuguese and American cuisine, with 138 and 16,056 recipes respectively. This heterogeneity suggests a substantial difference in the way ingredients are combined and possibly also in their frequency of use in recipes, consistently with previous findings on different data sets^{21,33}. The exponential growth of recipes with the number of ingredients (best-fit exponent $\gamma = 0.1$, $R^2 = 0.77$) allows a broad classification of cuisines into two groups—based on their position relative to the best-fitting curve. Cuisines above this curve, such as the Italian cuisine, feature more recipes than expected from their rough ingredient count (e.g., 378 ingredients forming 7493 recipes, surpassing the expected 4288). Conversely, cuisines below the curve indicate fewer recipes than those expected based on the number of ingredients employed. One example is the Portuguese cuisine, which has 144 ingredients but only 138 recipes.

The scaling of Fig. 1 reveals differences between cuisines at a macroscopic level, hiding preferences in the use of ingredients and on the way these are combined in recipes. To gain insights into these aspects, we examine the distribution of recipe sizes across cuisines. Figure 2a shows the probability distribution $P_r(n_r = n)$ that a randomly chosen recipe of a given cuisine has n ingredients. Each distribution represents a specific cuisine and is color-coded according to the cuisine’s continent affiliation (as depicted in Fig. 1). Cuisines are ordered from top to bottom based on their average recipe size. This analysis reveals two distinct tendencies: European (especially Continental and Southern European) cuisines tend to favor on average smaller recipe sizes, while Asian cuisines exhibit a preference for larger recipe sizes. Accordingly, the recipe with the largest number of ingredients in the data set is Vegetable Korma—a traditional Indian dish made with mixed vegetables, spices, and dairy products—which uses a staggering amount of 31 ingredients (see Supplementary Table S1 where we report the largest recipe of each cuisine, with the respective ingredients list).

A step further in highlighting the characteristics of cuisines is to focus on which specific ingredients are utilized and mixed in a recipe. In order to do this, we consider the division of ingredients into macro-categories or “ingredient types”. In our data set, ingredients are classified into 20 different types, namely: Additive, Bakery, Beverages, Alcoholic Beverages, Cereals, Dairy, Essential Oils, Fish, Flowers, Fruit, Fungi, Herbs, Legumes, Maize, Meat, Nuts & Seeds, Plants, Seafood, Spices, Vegetables. The rationale behind this classification, inherited from the original dataset⁴², is described in detail in the Methods section. We define the popularity of an ingredient type within a given cuisine as the fraction of ingredients of that type that constitute the recipes of that cuisine. Ingredient popularity enables us to shed light on how different cuisines utilize various ingredient categories within their recipes. To achieve this, in Fig. 2b we report a heatmap encoding the ingredient popularity profiles. Rows represent individual cuisines and columns represent specific ingredient types. To compute popularity profiles, we count the total number of occurrences of all ingredients belonging to a specific type across all recipes within a given cuisine. To account for potential biases arising from varying numbers of recipes or their lengths, we normalize these counts by dividing them by the sum of all recipe sizes within that cuisine. The heatmap illustrates a general pattern of ingredient usage across cuisines, with darker colors indicating higher popularity of certain ingredient types (on average, 9 types out of 20 are more popular than others). This common trend in ingredient usage is characterized by darker regions concentrated towards the left-hand side of the heatmap, indicating a higher popularity of certain ingredient types; as one moves towards the right-hand side, the popularity of ingredient types decreases and colors gradually lighten. However, despite this overall trend, individual cuisines exhibit distinct preferences. This becomes evident when individually examining the distribution of popularity profiles, as done for five specific cuisines in Fig. 2c (see Fig. S1 for the popularity profiles of the remaining cuisines). Furthermore, to emphasize differences among cuisines, we report in the heatmap of Fig. 2(d) the z-scores of the ingredient type popularity. Here, regions ranging from light-red to dark red in the heatmap, denote progressively higher positive z-scores, indicating that a cuisine tends to use certain ingredient types more frequently than others. Conversely, regions with light to dark blue colors represent increasingly negative z-scores, indicating below-average usage of certain ingredient types. For the sake of completeness, we report in Fig. 2e the z-scores distributions for five specific cuisines (see Fig. S2 for the z-scores distributions of the remaining cuisines). Notably, Fig. 2d, e reveal some distinctive traits of cuisines, which can be attributed to geographical, cultural or religious reasons. For example, India stands out as the cuisine that uses meat the least, with a high negative z-score, which reflects cultural restrictions on meat consumption. In contrast, Spices, Legumes and Flowers are used in Indian recipes with much higher frequency compared to other cuisines, aligning with the popular consensus of Indian cuisine as predominantly vegetarian, with extensive use of Flowers-based cooking oils and aromatic Spices. Similarly, the impact of geography on ingredient usage is exemplified by Scandinavian cuisine, which shows significantly lower usage of Vegetables, Herbs, and Plants compared to most other cuisines. The harsh climate of Scandinavian countries creates unfavorable conditions for the cultivation of most Vegetables, resulting in a distinct culinary approach that relies on other ingredients.

The networks of ingredient combinations

A cuisine is not solely characterized by the quantity and type of used ingredients, but also by the way in which the ingredients are combined. To describe the intricate patterns of ingredient combinations within recipes and better capture how such relationships vary across different cuisines, we introduce the network of ingredient combination. Recipe data sets can be represented as a bipartite graph, where ingredients and recipes are the two disjoint sets of nodes (see Methods for details). The projection of such networks on the set of ingredients (see Methods) allows us to obtain unipartite ingredient-ingredient graphs which provide summarized information on how ingredients are combined in recipes, while sensibly reducing the

number of nodes of the original graph. However, the projection graphs present two major obstacles: (i) the size reduction from bipartite to unipartite is not sufficient, leaving the resulting graphs overly complex for interpretation; (ii) the comparison between graphs might be inherently biased due to variations in graph size and ingredient sets across cuisines. To address these issues, we consider a coarse-grained representation of the ingredient-ingredient graphs, hereafter referred to as *ingredient type-ingredient type graphs*, or simply *type-type graphs*. Nodes in such graphs correspond to the 20 distinct ingredient types in our data set. An edge connects ingredient types t and t' if there is at least one recipe of the cuisine containing at least a pair of ingredients respectively of type t and t' . The edge weight reflects the frequency of such pairings across the entire data set. For example, the “apple strudel” recipe of Central European cuisine, comprising a total of 4 ingredients (apples, raisins, butter, puff pastry), ensures the existence of the 4 following edges: Fruit-Fruit, Fruit-Dairy, Fruit-Bakery, Bakery-Dairy. In particular, such recipe contributes with a weight of 2 to the two edges Fruit-Dairy and Fruit-Bakery in the network, reflecting the fact that butter and puff pastry are paired with two different fruits. The type-type graphs encode the main features of how ingredients are combined, without the complexity of a description at the level of individual ingredients. Moreover, it offers a uniform representation of culinary traditions across different cuisines. Being all made by the same number of nodes, our networks allow for direct comparisons of worldwide cuisines, revealing commonalities and discrepancies in ingredient interactions.

We constructed the ingredient type networks for each of the 23 cuisines. For better visualization of the obtained networks, due to their high density, we have used the disparity filter method introduced by Serrano et al.⁴³. With this method, we retain only the edges whose weights represent statistically significant deviations from a null model where the weights are uniformly distributed among the links of a node. Figure 3a–d show the filtered ingredient networks of 4 representative cuisines obtained by setting a p-value of 0.2 (see Fig. S3 for all other cuisines). We observe that the filtered networks of both American and Italian cuisines [Fig. 3a, c] preserve a high percentage of their original edges, while also exhibiting two very similar structures. Differently, Indian and Japanese filtered networks [Fig. 3b, d] have fewer edges. To gain deeper insights into how different cuisines combine food types, we analyze the z-scores of edge weights for specific type pairs across cuisines. For each type pair, we calculate the weight distribution across all cuisines and then derive the corresponding z-scores. In Fig. 3e, we report the scores for each cuisine using a strip plot. The resulting distributions of z-scores reveal interesting variations in the way cuisines use food type pairings. A cuisine with a long right tail in its z-score distribution indicates frequent use of certain type pairs compared to other cuisines. This trend is particularly noticeable in Asian cuisines, where less common ingredient pairings, such as the Flowers-Legumes combination in Indian cuisine or the Beverages-Fish pairing in Thai cuisine, are more prevalent. In contrast, cuisines like American, Canadian, and Australasian show a more uniform z-score distribution, with most scores falling within the -1 to $+1$ range. A likely explanation for this pattern is the history of mass immigration in these countries during the 20th century, which introduced a wide array of culinary traditions from different regions of the world. As these diverse influences blended, the distinctive characteristics of individual cuisines may have averaged out, leading to a more homogenized set of ingredient pairings.

To investigate the most important associations of food pairs for each of the 23 cuisines, we report in Table 1 the 5 edges with the largest weight, along with their corresponding weight fraction. Notably, for most cuisines, the top 5 edges (about 3% of the total number of edges) represent approximately 15% of the total weight. Additionally, the top 50 edges (about 25% of the total number of edges) capture about 75% of the total weight. The exception to this trend is observed in East Asian cuisines, where sensibly lower values are observed. The data in Table 1 reveals a noticeable recurrence of certain pairs of ingredient types within the top 5 pairs of each cuisine. For instance, pairings like *Spices – Vegetables* and *Additive-Vegetables* frequently appear among the top 5 in many cuisines, reflecting the prominence of the three

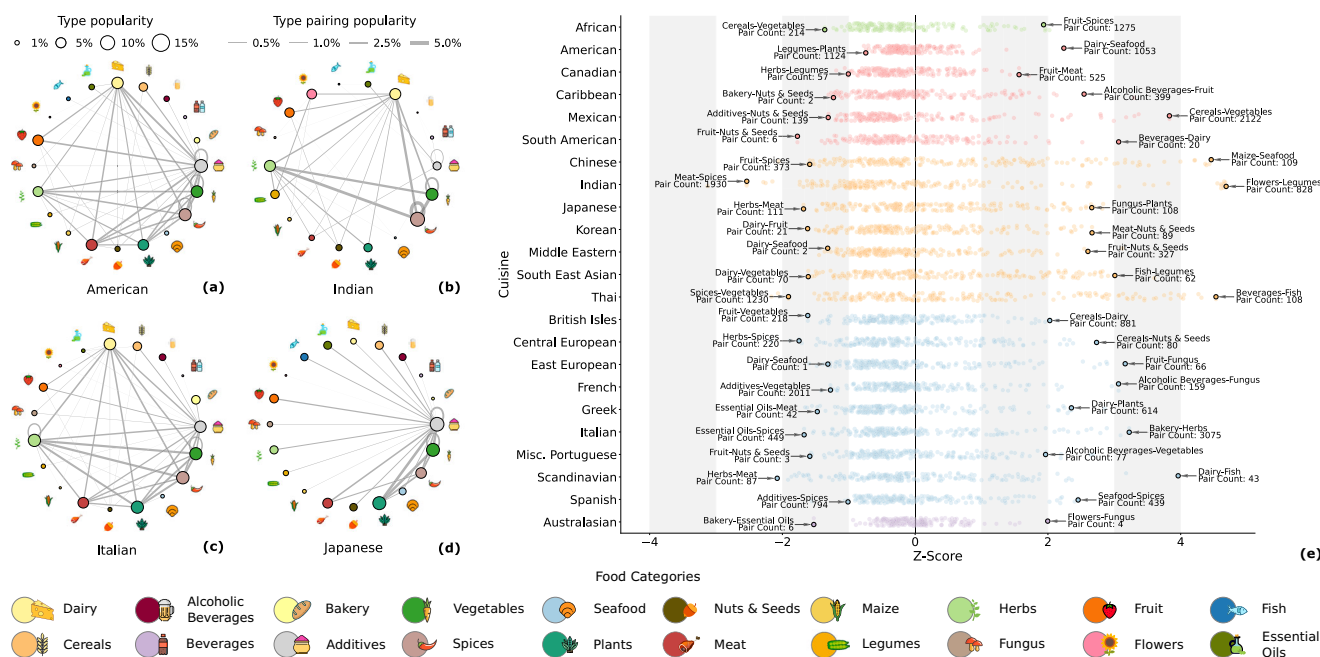


Fig. 3 | Networks of ingredient combinations and distributions of type pairing popularity across cuisines. Backbones of ingredient combinations networks for four selected cuisines: American (a), Indian (b), Italian (c) and Japanese (d). Nodes correspond to ingredient types and are drawn on a circular layout with additional type label symbols, with sizes proportional to the type’s recipe popularity in the cuisine. An edge is present in the network of ingredient combinations if two ingredient types occur in the same recipe at least once, and its thickness is proportional to the number of co-occurrences of types across all recipes. To highlight

differences between networks we show the network backbones obtained through the disparity filter⁴³, plotting only statistically significant edges with a p-value larger than 0.2 (e) Strip plots of z-scores for edge weights across world cuisines. For each cuisine on the y-axis, the x-axis shows the z-scores of edge weights relative to the distribution of the corresponding edge weights across all 23 cuisines. A jitter is applied for better visualization. Two edges are highlighted for each cuisine, one with a high z-score and one with a low z-score, representing significant deviations in the co-occurrence patterns.

most common ingredient types (Spices, Vegetables, and Additive) across cuisines. This observation is further supported by the analysis in Table 2 (left), where we report the 10 most frequent type pairings across all the 23 cuisines, i.e., the edges with the largest average weight. Here, Spices, Vegetables, and Additive constitute the first three edges in the ranking, underscoring their central importance across cuisines. Moreover, nine of the top ten pairs, except for the pair Additive - Meat, include either Spices or Vegetables. The recurrence of these type-type pairs across various cuisines suggests that certain ingredient combinations are particularly versatile and essential in the construction of diverse flavor profiles. This pattern implies that, while cuisines may differ in many aspects, they share common culinary strategies that rely on a core set of ingredient pairings. Table 2 (right) reports the top 10 type pairs ranked according to the standard deviation of weights across cuisines. In this list, the Dairy category is involved in four of the top five pairs. This is because dairy products are essentially a staple food for the cuisines of Central and Northern Europe, while being almost completely unknown to Asian culinary traditions.

To reduce the number of links in Fig. 3 and focus on meaningful associations between ingredients, we construct the maximum spanning trees (MSTs)²⁶ of the type-type graphs. Indeed, MSTs are well-suited for this task as they map a graph into a connected tree. The pruned graph is obtained minimizing the number of connections while maximizing the sum of the link weights. MSTs can be thus regarded as connected “minimal graphs”, easier to interpret and visualize, that give us the network fingerprint of a cuisine. The resulting MSTs for 6 of the 23 cuisines that we have studied are shown in Fig. 4 (see Fig. S4 of the Supplementary for all cuisines). Clear distinctions among these MSTs across cuisines underscore the effectiveness of this method in highlighting culinary differences otherwise not directly visible from the full networks. Notice that the obtained trees are all organized around one or two “central” nodes, i.e., the ones with the highest closeness centrality. For example, the Indian MST (Fig. 4c) appears as a star-like graph, with Spices at its core; this configuration effectively captures the

predominant centrality of Spices in Indian cuisine, which are an essential component to most of its recipes. These findings are in line with those of Jain et al.³⁴, who demonstrated the central role of Spices through the analysis of different regional cuisines of India. The negative food pairing tendency found across Indian cuisines, that is the repeated association of complementary ingredients sharing a low number of compounds, can indeed be attributed to the ubiquitous presence of Spices. We find a similar structure also in the MST describing the Mexican cuisine (Fig. 4e), where Spices are replaced by Vegetables. A different picture emerges instead from some other cuisines, that feature more than one ingredient type at their core. One example is the prominent role of cheese and dairy products in French cuisine (Fig. 4b) showing high affinity with certain ingredient types, like meat and mushrooms (*Fungi*), traditionally associated with umami flavor. The MST of the Italian cuisine (Fig. 4d) stands out for the central role played by Herbs (jointly with Spices), a unique feature shared only with Thai cuisine (see Fig. S4 of the Supplementary). A simple measure exemplifying structural differences across MSTs is their diameter, which quantifies how far apart the two most distant ingredient types are when moving along the shortest paths. In networks with small diameters, central ingredient types—typically those with the largest frequencies—tend to participate in a larger number of links, suggesting that these types are more evenly distributed across recipes compared to MSTs with larger diameters. Indian cuisine (Fig. 4c) presents the smallest diameter, corresponding for example to the 3 links connecting seafood to meat passing through Spices; among the MSTs shown here, American cuisine displays the largest diameter, corresponding for example to the 5 links connecting mushrooms to bakery via Vegetables, Spices, Additives, and Cheese. The highest diameter, formed by 6 links, is the one associated with Central European cuisine (see Fig. S4).

Identifying a cuisine from a set of recipes

We investigate here how well a world cuisine can be represented by its network of ingredient type combinations. We do this by setting up a series of

Table 1 | The second column reports the top 5 ingredient type-type pairs with the largest weight for each of the 23 cuisines

Cuisine	Top 5 Edges	Top 5 Weight Fraction	Top 50 Weight Fraction
African	- - - -	0.17	0.80
American	- - - -	0.12	0.74
Australasian	- - - -	0.13	0.73
British Isles	- - - -	0.16	0.78
Canadian	- - - -	0.14	0.78
Caribbean	- - - -	0.13	0.75
Central European	- - - -	0.15	0.78
Chinese	- - - -	0.11	0.65
East European	- - - -	0.15	0.80
French	- - - -	0.12	0.74
Greek	- - - -	0.15	0.82
Indian	- - - -	0.18	0.78
Italian	- - - -	0.14	0.76
Japanese	- - - -	0.15	0.68
Korean	- - - -	0.16	0.77
Mexican	- - - -	0.14	0.78
Middle Eastern	- - - -	0.14	0.78
Misc. Portuguese	- - - -	0.15	0.75
Scandinavian	- - - -	0.17	0.75
South American	- - - -	0.13	0.77
South East Asian	- - - -	0.12	0.68
Spanish	- - - -	0.15	0.71
Thai	- - - -	0.10	0.62

Additives

Bakery

Beverages

Alcoholic Beverages

Cereals

Dairy

Essential Oils

Fish

Flowers

Fruit

Fungus

Herbs

Legumes

Maize

Meat

Nuts & Seeds

Plants

Seafood

Spices

Vegetables

The third and the fourth columns report the total weight fraction accounted for by the top 5 and 50 type-type pairs with the largest weight, respectively.

classification experiments in which the task is to correctly associate a given set of recipes to the world cuisine from which they are extracted. The idea is to train a classification model on three different inputs of increasing complexity, namely the vector of ingredient (type) frequencies, the MST of ingredient (type) combinations, and the full network of ingredient (type) combinations.

In this way, it is possible to compare the predictive power of a cuisine representation based on the network of ingredient combinations, with respect to the other representations (ingredient frequencies, or MST of ingredient combinations).

For each of the three different representations, we independently train a Support Vector Machine (SVM) to predict each cuisine from a sample consisting of 15% of its recipes. More details on the bootstrapping procedure and experiment pipeline can be found in the Methods section. The SVMs

trained on the ingredient combination networks and their corresponding MSTs achieve an average classification accuracy of 0.95 (standard deviation 0.05) and 0.87 (standard deviation 0.06), respectively. Both graph-based representations outperform the SVM trained on ingredient frequencies, which only reaches an average accuracy of 0.79 (standard deviation 0.07). To get further insights on the classification performances beyond the accuracy scores, Fig. 5a–c report the average confusion matrices. We notice that models trained on ingredient frequencies mis-classify more cuisines than the other two models trained on ingredient combinations. In particular, the heatmap in Fig. 5a shows that cuisines belonging to either culturally or geographically close countries are often mixed, e.g., African with Middle Eastern, Greek with Italian, Korean with Japanese, and Portuguese with Spanish. By contrast, the darkest colors along the diagonals of the heatmaps in Fig. 5b, c show that, for both MST and network-based models, the most

Table 2 | (Left) Top 10 ingredient type-type pairs with the largest average weight

Type Pair	Average Weight	St. Dev.
Spices - Vegetables	0.027	0.006
Additive - Spices	0.025	0.004
Additive - Vegetables	0.023	0.005
Additive - Meat	0.022	0.007
Herbs - Vegetables	0.022	0.008
Vegetables - Vegetables	0.022	0.005
Herbs - Spices	0.021	0.008
Plants - Vegetables	0.021	0.006
Meat - Spices	0.020	0.004
Plants - Spices	0.020	0.006
Type Pair	St. Dev.	Average Weight
Additive - Dairy	0.013	0.020
Cereals - Dairy	0.012	0.014
Dairy - Meat	0.011	0.016
Additive - Cereals	0.010	0.017
Dairy - Spices	0.009	0.016
Additive - Additive	0.008	0.020
Herbs - Spices	0.008	0.021
Cereals - Meat	0.008	0.013
Herbs - Vegetables	0.008	0.022
Additive - Meat	0.007	0.022

Corresponding standard deviations are reported in the last column. **(Right)** Top 10 ingredient type-type pairs with the largest weight variation (standard deviation). Corresponding average weights are reported in the last column.

frequently predicted cuisine label always coincides with the true label. While a degree of ambiguity persists for some cuisines when using MSTs, this is almost entirely resolved with the use of the full network of ingredient (type) combinations. The only exception is the Misc. Portuguese cuisine, which is also the cuisine with the fewest number of recipes. Overall, the high classification accuracy achieved by the SVM based on the ingredient combination networks demonstrates their ability to efficiently encode rich information about cuisines. Their good performance suggests that networks of ingredient (type) combinations offer a more precise representation of a cuisine than the frequency of ingredients (type) alone. Moreover, it is worth stressing that even though MSTs are obtained by pruning 90% of the original edges, they are still able to achieve very high classification accuracy, further confirming the relevance of a networked approach.

Classification and clustering of world cuisines

We now exploit the networks of ingredient type combinations to uncover similarities and differences among 23 worldwide cuisines and to classify them into a smaller number of groups. We have adopted a hierarchical clustering based on the weighted adjacency matrices of the networks of ingredient types. We used the 'average' linkage method and Jensen-Shannon distance to quantify how "distant" cuisines are from each other. The full dendrogram is reported in Fig. 5d. Cutting at 0.21, the ingredient type graphs identify 6 main clusters, reported with different colors. Of these 6 groups, one uniquely identifies India, as proof that Indian cuisine has its unique way of combining ingredients in recipes. Mapping these culinary groups geographically, as illustrated in Fig. 5e, reveals three broader geographical clusters. Of these, one entirely coincides with Asian cuisines, further divided into two sub-clusters based on latitude (orange colors). Differently, the remaining clusters do not exhibit clear geographical separations (light green and light blue colors). However, the cluster including Central and Northern Europe (light blue color) encompasses Greek, English, and Italian cuisines, as well as what are referred to as "new world" countries. In a way, this cluster reflects the historical and cultural

influences shaped by colonialism in the evolution of world cuisines. This is also evident in the cluster that includes Spanish and Portuguese cuisines (light green), grouped with Latin American countries, formerly part of their colonial empires. Similarly, the presence of North Africa and the Middle East, easily explained by geographical continuity with the Iberian peninsula, could be the result of the long-lasting effects of medieval Islamic colonization of Southern Europe. Taken together, the overall clustering appears to depend on the interplay of geographical and cultural ties, which concur in shaping the ingredient profile of a cuisine.

Discussion

World cuisines embody a broad and intricate fusion of flavors, traditions, and cultural influences. Over time, this elaborate composition has garnered growing attention, not only due to its implications for health and dietary practices^{44,45}, but also for its role in inspiring new fields like computational gastronomy⁷. In this work, we have introduced a novel network-based approach that sheds light on both general trends and distinctive features of culinary traditions.

We first found that the number of recipes within a single cuisine exhibits an exponential relationship with the number of ingredients. Since the existing recipes constitute only a fraction of the potential ingredient combinations (as from other recipe repositories^{4,46}, only about 10^6 recipes exist out of an estimated $\sim 10^{15}$ potential combinations). This perceived limitation may embody a principle of least energy. It suggests that cuisines often rely on variations of successful dishes, instead of exploring completely novel combinations of ingredients. As highlighted by Kinouchi et al.²¹, modern cuisines undergo continual evolution, with high-fitness ingredients replacing lower-fitness ones. More recently, Nazari et al.⁴⁷ provided direct evidence supporting this hypothesis by reconstructing the evolutionary history of Italian stuffed pasta. Factors such as antimicrobial properties have been suggested to influence ingredient selection over time^{48,49}. Consequently, within a single cuisine, the space of potential ingredient combinations, even within the bounds of "gastronomical rules", remains only

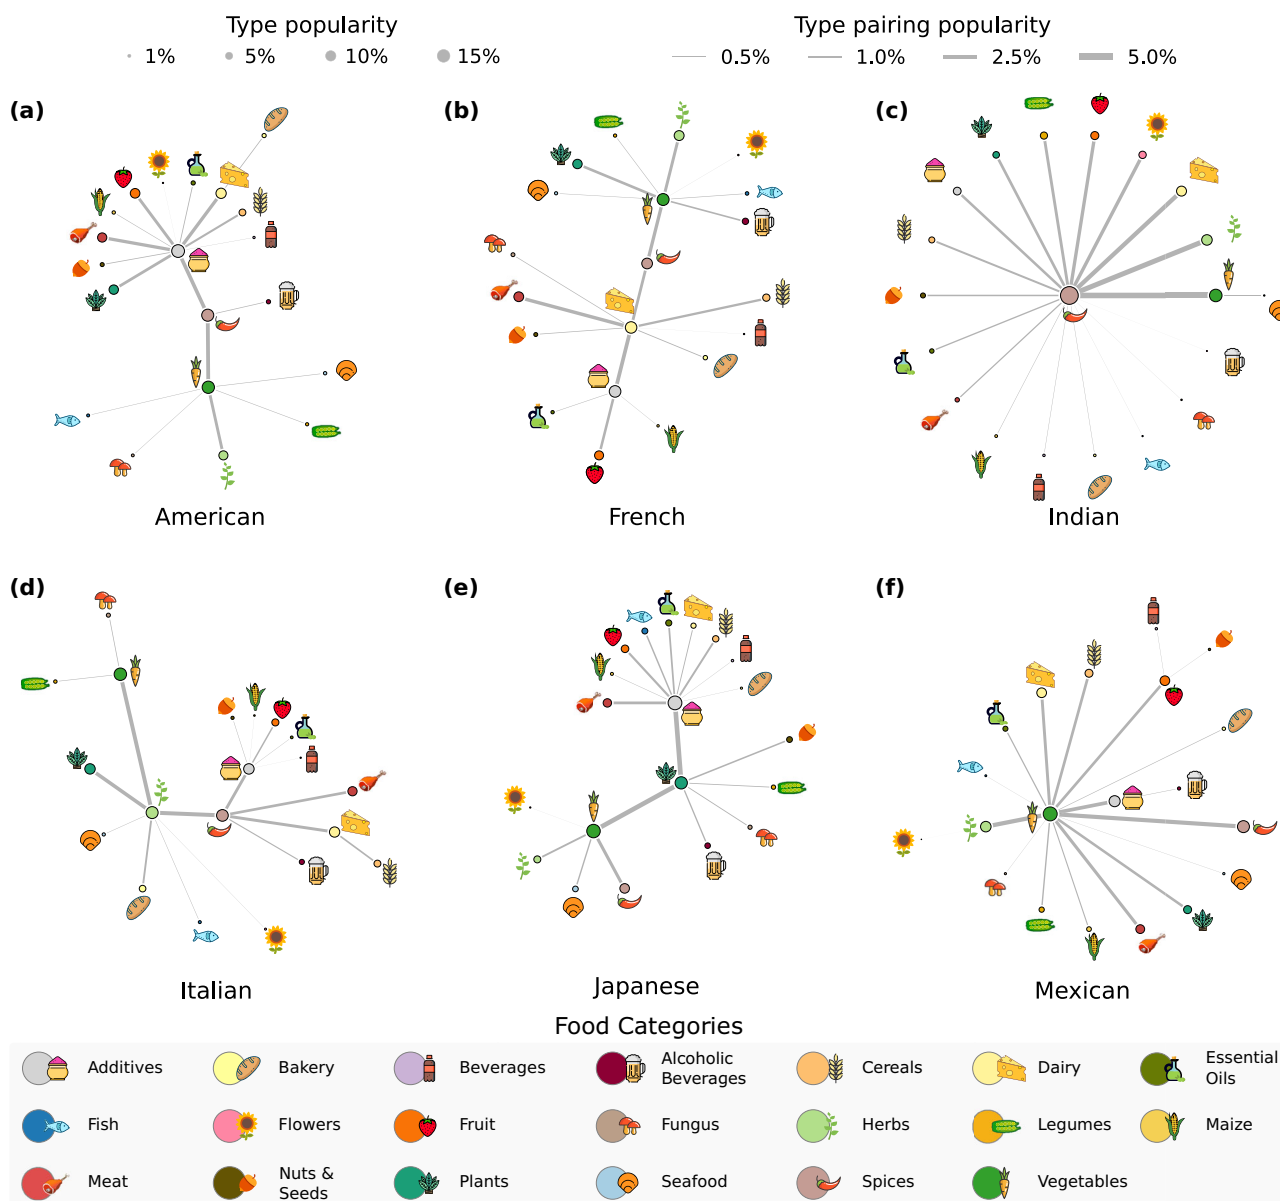


Fig. 4 | The backbone of ingredient combinations. Maximum spanning trees (MST) of the networks of ingredient combinations for six selected world cuisines: American (a), French (b), Indian (c), Italian (d), Japanese (e) and Mexican (f). The different topologies of the trees reveal the distinct overall organization of ingredient combinations. For instance, Indian cuisine exhibits a star-like topology centered

around Spices, whereas Italian and French cuisines display a more balanced reliance on different ingredient types, such as Dairy and Herbs. Node sizes are proportional to type popularity within the cuisine, while edge thickness corresponds to the weight in the original network of ingredient combinations, and indicates the number of co-occurrences of type pairs across all recipes.

partially explored. This ongoing evolutionary process contributes to the dynamic and adaptable nature of culinary traditions.

By examining how cuisines utilize ingredients of different types we also unveiled a pattern in how ingredient types are organized based on average prevalence. Cuisines generally tend to use the same ingredient types with comparable frequencies, suggesting a hierarchy. Categories like Vegetables, Spices, and Additive are extensively used across all cuisines, while others like Flowers, Beverages, and Maize exhibit overall low popularity. We hypothesize that various factors shape this pattern. For instance, the high prevalence of types like Vegetables may stem from their nutritional profile, providing key nutrients. Similarly, the prevalence of Additive and Spices can be justified by their specific culinary functions, such as flavoring or acting as chemical agents to catalyze certain reactions or provide antimicrobial properties to food. The networks of ingredient type combinations, introduced for the first time in this work, proved quite effective for classification tasks, showing that world cuisines can be essentially identified by how they

combine ingredients based on their types. Results presented in Fig. 5 show clearly that the unique and intertwined pattern of ingredient types constitutes a very good “fingerprinting” for each cuisine. This is also confirmed when considering dimensionality reduction methods, such as tSNE⁵⁰ (see Fig. S5 of the Supplementary for details), where ingredient type graphs form a culinary continuum that goes from East Asia to Scandinavia, where geographical proximity appears to be particularly relevant in determining the 2D projections. These findings align with those of Zhu et al.⁵¹, who, employing a distinct approach based on ingredient frequencies, analyzed regional cuisines in China and identified geographical proximity as the key factor of regional cuisine similarities. Furthermore, network measures derived from ingredient-type graphs (e.g., the spectral radius) show positive correlations with exogenous World Development Indicators, such as food imports and food insecurity, for the countries associated with that cuisine (see Fig. S6). One possible explanation is that cuisines with higher spectral radii might reflect a less varied combination of ingredient types, which could

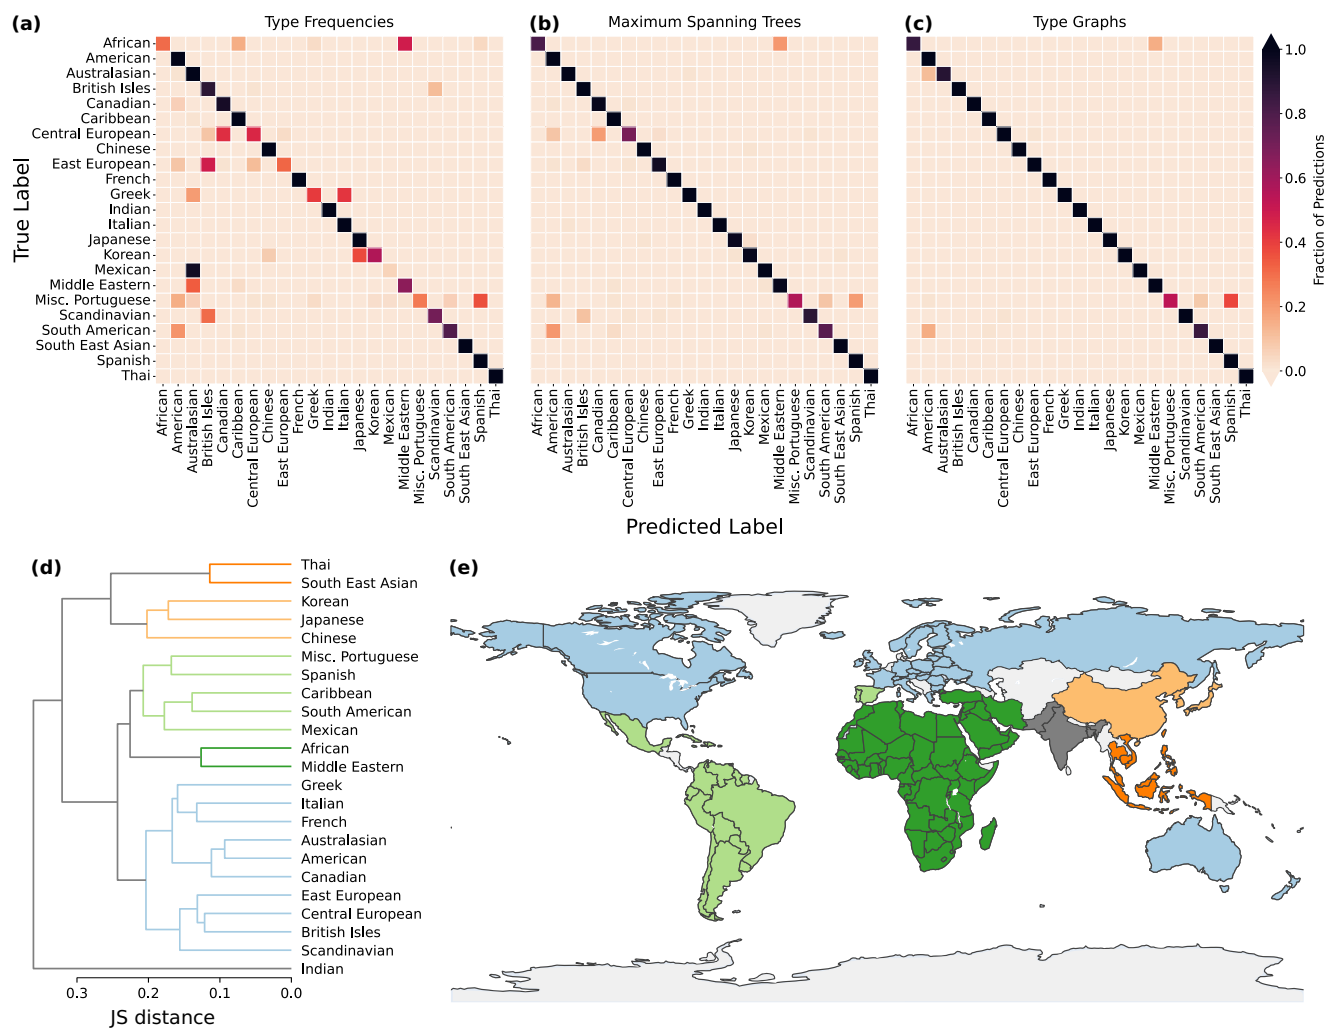


Fig. 5 | Identification and clustering of world cuisines. a–c Results of a Support Vector Machine (SVM) prediction model trained to associate a given set of recipes to the world cuisine they are extracted from. We report the confusion matrices obtained by using respectively: **a** ingredient frequency vectors, **b** Minimum spanning tree (MST) of ingredient combination networks, and **c** ingredient combination networks, as inputs to the SVM. **d**, **e** Hierarchical clustering of world cuisines based

on their ingredient combination networks. **d** Hierarchical tree obtained from an average linkage clustering method using a Jensen-Shannon (JS) distance. Cutting the tree at distance 0.21 produces six different groups, shown with different colors. **e** Geographical mapping of the resulting clusters from (d) highlighting the geographical distribution of the culinary groups.

correspond to more limited diets and thus correlate with higher food insecurity. Likewise, this measure may hint at shortages of certain ingredient types, potentially requiring greater reliance on imports. These preliminary observations highlight the need for further investigation to confirm and refine these relationships.

Finally, we showed how the analysis of the Maximum Spanning Trees (MSTs) of the graphs of ingredient type combinations, widely used to filter networks in finance^{52,53} and biology^{54,55}, can highlight the most relevant ingredient types and their pairings. Even though MSTs contain less information than the original type-type graphs, they provide a compact representation in terms of the essential components defining the culinary identity of a cuisine. In this respect, MSTs not only enable us to refine the visual representation of food type graphs by minimizing complexity and redundancy, but allow for a more intuitive exploration and analysis of the relationships between ingredients.

Since the structure of these graphs ultimately depends on how ingredients are categorized, it is important to stress the implications and limitations of our classification choices. We acknowledge that the ingredient classification we adopted in this work inherits intrinsic arbitrariness, carried over from the original dataset⁵⁶. First, the choice of categories may reflect potential cultural biases: ingredients that in our framework are grouped

under the same broad categories might be perceived as distinct food types in certain cultures. A notable example is that of seaweeds and algae, which are categorized under Seafood in our dataset based on their environmental (marine) origin. While this classification is consistent with ecological criteria and taxonomic ambiguity (algae are not strictly plants), it may overlook their cultural specificity, such as the central role of seaweeds in Japanese cuisine, potentially motivating their treatment as a separate ingredient category. This example highlights a broader point: culturally specific ingredient types offer a compelling reason to explore alternative classification systems, especially in regionally focused or cuisine-specific studies, where finer distinctions may lead to more meaningful insights. A second simplification in our framework concerns food processing: we merge different processed forms of the same raw item (e.g., fresh, dried, or canned tomatoes) into a single entity. Consequently, our categories capture the natural origin of ingredients but not the form they take after processing. While this choice allowed a first analysis by reducing complexity and assumptions, we acknowledge that processing can significantly influence ingredients' function and flavor, particularly in certain cuisines. A natural extension of our work would thus account for processing stages through a hierarchical categorization scheme, with subcategories (e.g., dried, fermented) nested within broader types (e.g., Vegetables, Dairy). Such more

granular taxonomy could enrich future analyses, if based on a clear and principled annotation of the database. Moreover, our analysis relies on data from a single source. Expanding the dataset to include a broader range of sources and cuisines would improve the robustness and generality of our findings. Additionally, our analysis considered a static snapshot of cuisines, using recipes available at a single point in time. However, against the well-established idea of national cuisines as part of the immutable traditions, culinary practices continuously evolve due to various factors such as technological advancements, ingredient availability, and local trends, all pushing for innovations in different directions. Hence, there is a need to recognize cuisines as dynamic and growing collections of recipes. Implementing a temporal network analysis would help to capture the evolution of culinary trends and ingredient popularity over time, and also to test the robustness of our network-based framework. Finally, any network approach primarily focuses on pairwise interactions, namely on the combination of pairs of ingredients. However, networks may not fully capture the complexity of culinary traditions. Extending our approach to higher-order networks, such as hypergraphs, to better describe higher-order interactions⁵⁷ among multiple ingredients, could provide a more nuanced understanding of ingredient combinations and their roles within recipes.

Methods

Dataset

CulinaryDB is a large repository consisting of 45772 recipes and two associated ingredient lists: one containing 929 individual ingredients, the other 103 compound ingredients, accounting for a total of 1032 unique ingredients. Of these, only 695 appear in the recipes themselves. The data were collected by CoSyLab⁴² from the following platforms: AllRecipes⁵⁸(16177), Food Network⁵⁹(15917), Epicurious⁶⁰(11069) and TarlaDalal⁶¹(2609). Recipes are grouped into 26 cuisines represented by world regions arching over 5 continents. A world region can either coincide with the national borders of a state or represent a broader geocultural area, such as Central America and Scandinavia (see Supplementary Note 1 and Table S2 for a detailed description of world regions). In the data set, each recipe is structured as a set of ingredients, which are classified into $T = 20$ different culinary *categories*, or *food types*, namely: Additive, Bakery, Beverages, Alcoholic Beverages, Cereals, Dairy, Essential Oils, Fish, Flowers, Fruit, Fungi, Herbs, Legumes, Maize, Meat, Nuts & Seeds, Plants, Seafood, Spices, Vegetables. Although this classification was defined by the creators of CulinaryDB⁵⁶, a comprehensive inspection shows that it follows a coherent set of principles grounded in biological origin, culinary function, and degree of processing. Ingredients are primarily categorized by their natural source, such as Meat (terrestrial animals), Fish (aquatic vertebrates), Seafood (aquatic invertebrates), Plants (leaves, stems, and roots), or Fungi, and secondarily, when applicable, by their dominant culinary role. For instance, aromatic leaves are labeled as Herbs, flavoring seeds as Spices, and concentrated extracts as Essential Oils. Items resulting from substantial transformation (e.g., Dairy, Bakery, Beverages) or used predominantly as process aids (e.g., Additives) are assigned to distinct classes. Borderline cases are resolved based on prevailing culinary identity: for example, honey and soy sauce, though processed, are categorized as Plants rather than Additives due to their natural origin and frequent culinary use. Manual inspection of all standardized ingredients confirms consistent application of these criteria.

For our analysis, we used data from 23 of the 26 cuisines, discarding those with fewer than 100 recipes to ensure statistical reliability, ending with a total of $R = 45661$ recipes. Additionally, we performed a preprocessing step to standardize and simplify the set of 695 ingredients found in recipes, resulting in a reduced set of $I = 604$ distinct ingredients. This included merging different shapes or varieties (e.g., all pasta types as ‘pasta’), unifying fresh, dried, or processed forms (e.g., dried and fresh peas as ‘peas’), and generalizing ingredient subtypes (e.g., cherry, plum, and globe tomatoes as ‘tomatoes’). We also aggregated ingredients produced through similar transformative processes into broader categories (e.g., all cheeses as ‘cheese’, all grape-based fermented drinks as ‘wine’), and grouped different animal parts under their corresponding species (e.g., chicken wings, thighs, and

liver as ‘chicken’). A detailed description of the preprocessing rules and illustrative examples is provided in Supplementary Note 1.

For each cuisine c , with $c = 1, \dots, 23$, we denote as \mathcal{R}_c the set of recipes and, as $R_c = |\mathcal{R}_c|$, the number of recipes in cuisine c , representing the size of the cuisine. Each recipe $r \in \mathcal{R}_c$ can be viewed as a subset of the finite set of ingredients \mathcal{I}_c used in cuisine c , from which dishes are created through different combinations. We indicate the number of ingredients in recipe r as n_r , and the number of ingredients in cuisine c as $I_c = |\mathcal{I}_c|$, with $I_c \leq I$. We indicate an ingredient type as t , with $t = 1, \dots, T$, and as t_i the type of an ingredient i , with $i = 1, 2, \dots, I_c$. We define the popularity $f_c(t)$ of an ingredient type t in cuisine c as the fraction of ingredients of type t in the recipes of cuisine c . To calculate $f_c(t)$, we sum the number of ingredients of type t across all the recipes in cuisine c , and normalize this sum by the total number of ingredients in all recipes of cuisine c . Similarly, we define the popularity $p_c(t)$ of an ingredient type t across the recipes of cuisine c , as the fraction of recipes of c containing at least an ingredient of type t . In other words, $f_c(t)$ represents the probability that a randomly extracted ingredient from recipes of cuisine c belongs to type t . Differently, $p_c(t)$ is the probability that a randomly extracted recipe of cuisine c contains at least an ingredient of type t . Formally, these definitions can be expressed as follows:

$$f_c(t) = \frac{\sum_{r \in \mathcal{R}_c} \sum_{i \in r} \delta(t_i, t)}{\sum_{r \in \mathcal{R}_c} n_r}, \quad (1)$$

and

$$p_c(t) = \frac{\sum_{r \in \mathcal{R}_c} \mathbf{1}(\exists i \in r | t_i = t)}{R_c} \quad (2)$$

where $\delta(a, b)$ is the Kronecker delta function, which is 1 if $a = b$ and 0 otherwise, and $\mathbf{1}(\bullet)$ is the indicator function that is 1 if any ingredient in the recipe belongs to type t and 0 otherwise.

Network representations of the data

We represent the data using three different network structures. The first is the *Recipe-Ingredient graph*, which is a bipartite graph²⁵ representation of the recipes of a cuisine c . The node sets of the graph consist of recipes and ingredients appearing in cuisine c . Specifically, if an ingredient i is part of a recipe r , an undirected edge connects i to r . The obtained graph is thus bipartite, as there are no edges inside the two separate sets of ingredients and recipes. It is straightforward to show that the degree distribution $P_r(k)$ for the recipes coincides with the recipe size distribution $P_r(n)$.

The second representation is the *Ingredient-Ingredient graph*, defined as the weighted projection⁶² of the recipe-ingredient graph on the set of ingredients \mathcal{I}_c . A pair of ingredient nodes i and j is connected if there exists at least one recipe containing both i and j . The weight w_{ij} is equal to the number of recipes of the cuisine that uses both i and j as ingredients.

The third representation is based on ingredient type graphs. To construct the type graph for a cuisine c we employ the following procedure. For each recipe r within the cuisine, we construct the multi-set⁶³ of type pairs t_i, t_j for every unordered pair of ingredients i, j . We then compute the co-occurrence frequency $n(t_i, t_j)$ of each type pair t_i, t_j across the multi-sets. We add the undirected edge (t_i, t_j) to the network, with the edge weight w_{t_i, t_j} set equal to the computed co-occurrence frequency. We finally enrich the network nodes with relative abundance (v_t) of the ingredient types, computed as the fraction of recipes that contain at least one ingredient of type t . Figure 3a–d shows the backbones of the resulting food type networks for 4 of the 23 worldwide cuisines. Node sizes in the plots represent relative frequencies (v_t) of different food types, with larger nodes indicating more prevalent food categories within a cuisine.

See Supplementary Note 1 and Fig. S7 for a graphical representation.

Classification experiment

The classification experiment comprises 250 repeated classification tasks, each involving a different data set comprising 23 classes corresponding to

the different cuisines. Each data point in our dataset is represented as a pair (X, y) , where X denotes one of the three possible feature representations of a recipe (type frequency vectors, ingredient type graphs, and their respective MSTs), and y is the corresponding cuisine label.

For each cuisine, we reserve 85% of the available recipes for training and leave the remaining 15% for testing. To create the training set, we follow a systematic sampling approach to ensure a balanced representation across all classes. For each of the 23 cuisines, we sample 75 recipes from the 85% reserved for training and use these to generate the feature representation X , obtaining a data point (X, y) for each of the three feature representations, with y being the associated class (cuisine) label. We repeat this process 100 times for each class (cuisine), resulting in three separate balanced training sets, each containing 2,300 data points (100 samples \times 23 classes). For the test set, we use the remaining 15% of recipes to generate a single data point (X, y) for each cuisine, resulting in a minimal balanced test set of size 23.

We utilize a Support Vector Machine (SVM) classifier for the classification tasks. For each of the three feature representations, a distinct SVM model is trained on the corresponding balanced training set. The performance of each SVM model is then evaluated on the respective test set, which contains one data point for each of the 23 cuisines.

To ensure the robustness and consistency of our results, the entire process—data preparation, training set construction, model training, and evaluation is repeated 250 times. This approach allows us to account for variability in data sampling and provides a comprehensive assessment of the classifier's performance across different feature representations and cuisine classes.

The performance associated with each feature representation is quantified in terms of the average test accuracy of the model trained using that representation across all 250 iterations. We use the standard deviation of the test accuracy as a measure of the error. This approach allows us to evaluate the stability and reliability of the classifier's performance across multiple iterations and feature sets.

Data availability

The data used in this manuscript are available from the original source (<https://cosylab.iitd.edu.in/culinarydb/>). All scripts supporting the findings of this study can be found at <https://github.com/claudiocaprioli/ingredient-combination>.

Code availability

All scripts supporting the findings of this study can be found at <https://github.com/claudiocaprioli/ingredient-combination>.

Received: 30 January 2025; Accepted: 23 September 2025;

Published online: 20 November 2025

References

- Bottéro, J. The culinary tablets at Yale. *J. Am. Orient. Soc.* **107**, 11–19 (1987).
- Nadeau, R. Apicius: a critical edition with an introduction and an English translation of the Latin recipe text by Christopher Grocock and Sally Grainger. *Food Cult. Soc.* **12**, 237–241 (2009).
- Zohar, I. et al. Evidence for the cooking of fish 780,000 years ago at Geshert Benot Ya'akov, Israel. *Nat. Ecol. Evol.* **6**, 2016–2028 (2022).
- Pilcher, J. M. *The Oxford Handbook of Food History*. <https://doi.org/10.1093/oxfordhb/9780199729937.001.0001> (Oxford University Press, 2012).
- Diamond, J. *Guns, Germs, and Steel: The Fates of Human Societies*. Book club kit. https://books.google.it/books?id=kLKTa_OeoNIC (W.W. Norton, 1997).
- Civitello, L. *Cuisine and Culture: A History of Food & People*. <https://books.google.it/books?id=HgiHAAAAYAAJ> (Wiley, 2004).
- Goel, M. & Bagler, G. Computational gastronomy: a data science approach to food. *J. Biosci.* **47**, 12 (2022).
- Goel, M. et al. Ratatouille: A tool for novel recipe generation. In *2022 IEEE 38th International Conference on Data Engineering Workshops (ICDEW)*, 107–110 (IEEE, 2022).
- Yamakata, Y., Imahori, S., Sugiyama, Y., Mori, S. & Tanaka, K. Feature extraction and summarization of recipes using flow graph. In *Social Informatics* (eds Jatowt, A. et al.), 241–254 (Springer International Publishing, 2013).
- Mori, S., Sasada, T., Yamakata, Y. & Yoshino, K. A machine learning approach to recipe text processing. In *Proc. 1st Cooking with Computer Workshop*, 29–34 (Citeseer, 2012).
- Wu, X. et al. A large-scale benchmark for food image segmentation. In *Proc. 29th ACM International Conference on Multimedia*, MM '21, 506–515. (Association for Computing Machinery, 2021).
- Khurana, D., Koli, A., Khatter, K. & Singh, S. Natural language processing: state of the art, current trends and challenges. *Multimed. Tools Appl.* **82**, 3713–3744 (2023).
- Bishop, C. *Pattern Recognition and Machine Learning: All "Just the Facts 101" Material*. Information science and statistics. <https://books.google.it/books?id=HL4HrgEACAAJ> (Springer, 2013).
- Marin, J. et al. Recipe1m+: a dataset for learning cross-modal embeddings for cooking recipes and food images. In *IEEE Trans. Pattern Anal. Mach. Intell.* (IEEE, 2019).
- Salvador, A., Drozdal, M., Giró-i Nieto, X. & Romero, A. Inverse cooking: Recipe generation from food images. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10453–10462 (IEEE, 2019).
- Channam, A., Swarup, B. R. & Rao, S. G. Extraction of recipes from food images by using CNN algorithm. In *2021 Fifth International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC)*, 1308–1315 (IEEE, 2021).
- Menichetti, G. & Barabási, A.-L. Nutrient concentrations in food display universal behaviour. *Nat. Food* **3**, 375–382 (2022).
- Menichetti, G., Ravandi, B., Mozaffarian, D. & Barabási, A.-L. Machine learning prediction of the degree of food processing. *Nat. Commun.* **14**, 2312 (2023).
- Aiello, L. M., Schifanella, R., Quercia, D. & Del Prete, L. Large-scale and high-resolution analysis of food purchases and health outcomes. *EPJ Data Sci.* **8**, 1–22 (2019).
- Wang, X.-W. et al. Nutritional redundancy in the human diet and its application in phenotype association studies. *Nat. Commun.* **14**, 4316 (2023).
- Kinouchi, O., Diez-García, R. W., Holanda, A. J., Zambianchi, P. & Roque, A. C. The non-equilibrium nature of culinary evolution. *N. J. Phys.* **10**, 073020 (2008).
- Kauffman, S. A. Investigations: the nature of autonomous agents and the worlds they mutually create. In *SFI working papers* (Santa Fe Institute, 1996).
- Tria, F., Loreto, V., Servedio, V. D. P. & Strogatz, S. H. The dynamics of correlated novelties. *Sci. Rep.* **4**, 1–8 (2014).
- Iacopini, I., Milojević, S. & Latora, V. Network dynamics of innovation processes. *Phys. Rev. Lett.* **120**, 048301 (2018).
- Newman, M., Barabási, A. & Watts, D. *The Structure and Dynamics of Networks*. Princeton Studies in Complexity. <https://books.google.it/books?id=6LvQIIP0TQ8C> (Princeton University Press, 2006).
- Latora, V., Nicosia, V. & Russo, G. *Complex Networks: Principles, Methods and Applications* (Cambridge University Press, 2017).
- Ahnert, S. E. Network analysis and data mining in food science: the emergence of computational gastronomy. *Flavour* **2**, 4 (2013).
- Bellingeri, M. et al. The recipe similarity network: a new powerful algorithm to extract relevant information from cookbooks (2025).
- Guillaume, J.-L. & Latapy, M. Bipartite structure of all complex networks. *Inf. Process. Lett.* **90**, 215–221 (2004).
- Guillaume, J.-L. & Latapy, M. Bipartite graphs as models of complex networks. *Phys. A Stat. Mech. Appl.* **371**, 795–813 (2006).

31. Torres, L., Blevins, A. S., Bassett, D. & Eliassi-Rad, T. The why, how, and when of representations for complex systems. *SIAM Rev.* **63**, 435–485 (2021).
32. Blumenthal, H. *The big fat duck cookbook* (Bloomsbury, 2008).
33. Ahn, Y.-Y., Ahnert, S. E., Bagrow, J. P. & Barabási, A.-L. Flavor network and the principles of food pairing. *Sci. Rep.* **1**, 196 (2011).
34. Jain, A. et al. Spices form the basis of food pairing in indian cuisine. *Phys. Soc.* <https://api.semanticscholar.org/CorpusID:17788890> (2015).
35. Jain, A., N K, R. & Bagler, G. Analysis of food pairing in regional cuisines of india. *PLoS One* **10**, e0139539 (2015).
36. Bagler, G. & Singh, N. Data-driven investigations of culinary patterns in traditional recipes across the world. In *2018 IEEE 34th International Conference on Data Engineering Workshops (ICDEW)*, 157–162 (IEEE, 2018).
37. Simas, T., Ficek, M., Diaz-Guilera, A., Obrador, P. & Rodriguez, P. R. Food-bridging: a new network construction to unveil the principles of cooking. *Front. ICT* **4**. <https://www.frontiersin.org/articles/10.3389/fict.2017.00014> (2017).
38. Barabási, A.-L., Menichetti, G. & Loscalzo, J. The unmapped chemical complexity of our diet. *Nat. Food* **1**, 33–37 (2020).
39. Kim, S., Sung, J., Foo, M., Jin, Y.-S. & Kim, P.-J. Uncovering the nutritional landscape of food. *PLoS One* **10**, e0118697 (2015).
40. Cenikj, G. et al. From language models to large-scale food and biomedical knowledge graphs. *Sci. Rep.* **13**, 7815 (2023).
41. Menichetti, G., Barabási, A.-L. & Loscalzo, J. Chemical complexity of food and implications for therapeutics. *N. Engl. J. Med.* **392**, 1836–1845 (2025).
42. CulinaryDB. <https://cosylab.iitd.edu.in/culinarydb/> (Accessed Jan 2021).
43. Ángeles Serrano, M., Boguñá, M. & Vespignani, A. Extracting the multiscale backbone of complex weighted networks. *Proc. Natl. Acad. Sci. USA* **106**, 6483–6488 (2009).
44. Cena, H. & Calder, P. C. Defining a healthy diet: evidence for the role of contemporary dietary patterns in health and disease. *Nutrients* **12**, 334 (2020).
45. Vajdi, M. & Farhangi, M. A. A systematic review of the association between dietary patterns and health-related quality of life. *Health Qual. Life Outcomes* **18**, 337 (2020).
46. Bieñ, M. et al. RecipeNLG: A cooking recipes dataset for semi-structured text generation. In *Proc. 13th International Conference on Natural Language Generation*, 22–28. (Association for Computational Linguistics, Dublin, Ireland, 2020).
47. Nazari, V. et al. Evolution of the italian pasta ripiena: the first steps toward a scientific classification. *Discov. Food* **4**, 57 (2024).
48. Billing, J. & Sherman, P. W. Antimicrobial functions of spices: why some like it hot. *Q. Rev. Biol.* **73**, 3–49 (1998).
49. Sherman, P. W. & Hash, G. A. Why vegetable recipes are not very spicy. *Evol. Hum. Behav.* **22**, 147–163 (2001).
50. van der Maaten, L. & Hinton, G. Visualizing data using t-sne. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).
51. Zhu, Y.-X. et al. Geography and similarity of regional cuisines in china. *PLoS One* **8**, 1–8 (2013).
52. Mantegna, R. Hierarchical structure in financial markets. *Eur. Phys. J. B* **11**, 193–197 (1999).
53. Valle, M. A., Ruz, G. A. & Morras, R. Market basket analysis: complementing association rules with minimum spanning trees. *Expert Syst. Appl.* **97**, 146–162 (2018).
54. King, B. M. & Tidor, B. MIST: maximum information spanning trees for dimension reduction of biological data sets. *Bioinformatics* **25**, 1165–1172 (2009).
55. Wu, Y., Bhat, P. R., Close, T. J. & Lonardi, S. Efficient and accurate construction of genetic linkage maps from the minimum spanning tree of a graph. *PLoS Genet.* **4**, e1000212 (2008).
56. Complex Systems Laboratory (CoSyLab), IIT-Delhi. Culinarydb: Data analytics for world cuisines. <https://cosylab.iitd.edu.in/culinarydb/> (2017). (Accessed Jan 2021).
57. Battiston, F. et al. Networks beyond pairwise interactions: Structure and dynamics. *Phys. Rep.* **874**, 1–92 (2020). Networks beyond pairwise interactions: Structure and dynamics.
58. Allrecipes. Allrecipes | recipes, how-tos, videos and more. <https://www.allrecipes.com/> (Accessed Jan 2021).
59. Food Network. Food network: Easy recipes, healthy eating ideas and chef recipes. <https://www.foodnetwork.com/> (Accessed Jan 2021).
60. Epicurious. Epicurious – recipes, menu ideas, videos & cooking tips. <https://www.epicurious.com/> (Accessed Jan 2021).
61. TarlaDalal.com. Tarla dalal – indian recipes. <https://www.tarladalal.com/> (Accessed Jan 2021).
62. Barrat, A., Barthélemy, M., Pastor-Satorras, R. & Vespignani, A. The architecture of complex weighted networks. *Proc. Natl. Acad. Sci.* **101**, 3747–3752 (2004).
63. Hein, J. *Discrete Mathematics*. Discrete Mathematics and Logic Series (Jones and Bartlett Publishers, 2003). <https://books.google.it/books?id=qLNfpb1hBWwC>.

Acknowledgements

A.S. acknowledges support from the SNSF COST project (grant no. IZCOZ0_198144).

Author contributions

All authors designed the study. C.C. and S.K. analysed the data. All authors interpreted the results, wrote and approved the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41538-025-00588-4>.

Correspondence and requests for materials should be addressed to Iacopo Iacopini or Andrea Santoro.

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) 2025