

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

A Survey on Hypergraph Representation Learning

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1916410> since 2023-12-28T16:11:34Z

Published version:

DOI:10.1145/3605776

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

A Survey on Hypergraph Representation Learning

ALESSIA ANTELM^{*}, Università degli Studi di Torino, Italy

GENNARO CORDASCO, Università della Campania “Luigi Vanvitelli”, Italy

MIRKO POLATO^{*}, Università degli Studi di Torino, Italy

VITTORIO SCARANO, Università degli Studi di Salerno, Italy

CARMINE SPAGNUOLO, Università degli Studi di Salerno, Italy

DINGQI YANG, University of Macau, Macau SAR, China, China

Hypergraphs have attracted increasing attention in recent years thanks to their flexibility in naturally modeling a broad range of systems where high-order relationships exist among their interacting parts. This survey reviews the newly born hypergraph representation learning problem, whose goal is to learn a function to project objects - most commonly nodes - of an input hyper-network into a latent space such that both the structural and relational properties of the network can be encoded and preserved. We provide a thorough overview of existing literature and offer a new taxonomy of hypergraph embedding methods by identifying three main families of techniques, i.e., spectral, proximity-preserving, and (deep) neural networks. For each family, we describe its characteristics and our insights in a single yet flexible framework and then discuss the peculiarities of individual methods, as well as their pros and cons. We then review the main tasks, datasets, and settings in which hypergraph embeddings are typically used. We finally identify and discuss open challenges that would inspire further research in this field.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Mathematics of computing** → **Hypergraphs**; • **Computing methodologies** → **Machine learning**.

Additional Key Words and Phrases: Hypergraph Representation Learning, Hypergraph Embedding, Hypergraph Neural Networks, Hypergraph Convolution, Hypergraph Attention

ACM Reference Format:

Alessia Antelmi, Gennaro Cordasco, Mirko Polato, Vittorio Scarano, Carmine Spagnuolo, and Dingqi Yang. 2022. A Survey on Hypergraph Representation Learning. 1, 1 (June 2022), 45 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Hypergraphs are the natural representation of a broad range of systems where group (or high-order or many-to-many) relationships exist among their interacting parts. Technically speaking, a hypergraph is a generalization of a graph where a (hyper)edge allows the connection of an arbitrary number of nodes [28]. Such structures can easily abstract social systems where individuals interact in groups of any size [17, 117]; for instance, in the case of a co-authorship

^{*}Corresponding authors.

Authors' addresses: [Alessia Antelmi](mailto:alessia.antelmi@unito.it), alessia.antelmi@unito.it, Università degli Studi di Torino, Torino, Italy; [Gennaro Cordasco](mailto:gennaro.cordasco@unicampania.it), Università della Campania “Luigi Vanvitelli”, Caserta, Italy, gennaro.cordasco@unicampania.it; [Mirko Polato](mailto:mirko.polato@unito.it), Università degli Studi di Torino, Torino, Italy, mirko.polato@unito.it; [Vittorio Scarano](mailto:vitsca@unisa.it), Università degli Studi di Salerno, Fisciano, Italy, vitsca@unisa.it; [Carmine Spagnuolo](mailto:cpagnuolo@unisa.it), Università degli Studi di Salerno, Fisciano, Italy, cpagnuolo@unisa.it; [Dingqi Yang](mailto:dingqiyang@um.edu.mo), University of Macau, Macau SAR, China, Macau SAR, China, dingqiyang@um.edu.mo.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

1




Fig. 1. An example of co-authorship hypergraph where each node represents an author, and each hyperedge connects all authors that have contributed to the same article.

collaboration network, a hyperedge may represent an article and link together all authors (nodes) having collaborated on it [10, 55, 57, 85, 97, 138, 160, 164, 182, 195] (see Figure 1). Similar situations, characterized by high-order interactions, also exist in biology [65, 151], ecology [54, 67], and neuroscience [94, 121]. Despite their powerful expressiveness, hypergraphs have been underexplored in literature (in favor of their graph counterpart) because of their inherent complexity and the lack of appropriate tools and algorithms. Recently, the trend has been drifting, thanks to a rising number of systematic studies demonstrating how the transformation of a hypergraph to a classical graph either leads to an inevitable loss of information or creates a large number of extra nodes/edges that increases space and time requirements in downstream graph analytic tasks [2, 17, 199, 203]. Specifically, hypergraphs have been proven to be a critical tool to use when the underlying system to study exhibits highly non-linear interactions between its constituents [17]. In practice, this consideration translates into using hyperedges to model (possibly indecomposable) group interactions that cannot be described simply in terms of dyads (and, hence, via graphs). For instance, hypergraph modeling has been exploited to embed key sociological concepts such as homophily (i.e., the influence a group exerts on a single individual) and conformity (i.e., group pressure, namely the tendency of an individual to align their beliefs to those of their peers, often reinforced by the nature of shared opinions) to investigate the dynamics of opinion formation [92, 131, 132, 149] and social influence diffusion [6, 163, 180, 222, 223] when groups are explicitly taken into account. Similarly, hypergraphs have been used to model epidemic-spreading processes to expressly account for community structure and non-linear infection pressure [5, 25, 82], and group dynamics [45, 93, 106, 122]. Throughout this review, we will discuss other application scenarios in which hypergraph modeling could be more beneficial than traditional graph modeling and elaborate on the characteristics of the many-to-many relations abstracted and the specific tasks addressed.

All graph-related problems and corresponding challenges still hold for the hypergraph-based setting, where the computational cost is even more significant due to the presence of high-order interactions [127, 155]. In this sense, the task of hypergraph representation learning (a.k.a. hypergraph embedding) further assumes a critical role in effectively and efficiently solving analytic problems. Embedding a network - either a graph or a hypergraph - means projecting its structure and possibly additional information onto a low-dimensional space where the structural and semantic information (e.g., nodes' neighborhood and features) is ideally preserved. The underlying idea of this procedure is that representing the nodes and (hyper)edges as a set of low dimensional vectors allow the efficient execution of the traditional vector-based machine learning algorithms on the (hyper)graph. As for graphs, the problem of hypergraph embedding, thus, lies in the overlap of two traditional research problems: hypergraph analytics and representation learning [30].




Fig. 2. Number of peer-reviewed publications related to hypergraphs from 1970 to 2021 (source: Scopus).

While the first problem aims to mine useful information from the hyper-network structure, representation learning seeks to learn compact representations (i.e., latent feature vectors) when addressing, for instance, classification [10, 159, 221], link prediction [52, 200, 217], and recommendation [175, 176, 206] tasks. Learning a latent representation of hypergraphs, rather than graphs, enables the exploration of the high-order correlation among data and the indecomposable nature of certain group relations to build more comprehensive representations, leading to better performance in practice [156, 204].

As witnessed in Figure 2, for the last decade, a growing body of work has devoted its efforts to investigating hypergraphs to design more effective solutions in various domains [17]. Still, so far, there is no systematic exploration of hypergraph embedding methods, as discussed in the following. This survey aims to fill this gap by providing a thorough overview of existing literature and offering a taxonomy of hypergraph embedding techniques. Our broadest intent is to develop a comprehensive understanding and critical assessment of the knowledge of this newly born research area.

Differences with previous surveys. Current surveys [63, 212] closely related to ours substantially differ in the topics discussed and the literature covered.

The work from Gao et al. [63] deals with the hypergraph learning problem (sometimes called hypergraph regularization), which is a related but different topic than hypergraph representation learning. According to [63], hypergraph learning is the process of passing information along the hypergraph topology in analyzing the structured data and solving problems such as node classification. Learning hypergraph embeddings is not the goal of hypergraph learning, although both tasks share some concepts and ideas (e.g., [221]). Further, the authors thoroughly analyze hypergraph generation methods (not covered in this survey).

The most recent survey from Zhang et al. [212] presents a shallow and brief excursus on hypergraph representation learning techniques, offering a single-level taxonomy similar to the one provided in this survey (see Section 5). However, the authors limit the discussion to a few representative works, focusing on how to handle uncertain data using hypergraphs. Their survey also covers some graph representation learning and hypergraph generation methods.

For the above reasons, our work can be regarded as a complement to the surveys by Gao et al. [63] and Zhang et al. [212] since it emphasizes the task of learning hypergraph embedding (a.k.a. hypergraph representation learning), providing a series of novel contributions that are listed below.

Contributions. Our contributions can be summarized as follows:

- *Inherent challenges.* As hypergraphs are a generalization of graphs, some challenges are directly inherited by the graph representation learning problem. However, the high-order nature of hypergraphs imposes additional

difficulties. We discuss the classical challenges of (hyper)graph embedding to then detail the peculiar challenges these structures pose.

- *New taxonomies.* We propose three taxonomies, classifying hypergraph embedding methods based on (i) their learning approach (spectral, proximity-preserving, and neural network techniques), (ii) the structure of the input hypergraph (homogeneous/heterogeneous, undirected/directed, uniform/non-uniform, static/dynamic, attribute/not attributed nodes, transformation into a graph) and (iii) the desired output (node/hyperedge embedding).
- *Comprehensive review.* The recent booming of the hypergraph representation learning field enabled us to collect, systematically review and characterize the whole evolution of this research area. We describe the most representative approaches of each hypergraph embedding class, unraveling the pros and cons of each learning mechanism and highlighting the high-level connections among the described techniques within each category.
- *Hypergraph analytic tasks.* We review the discussed methods under the lens of representation learning applications, categorizing them according to node-related and hyperedge-related tasks.
- *Future directions.* We examine the limitations of the current state-of-the-art and propose six potential future directions in the area in terms of problem setting, modeling techniques, interpretability, and scalability.

Article organization. The remainder of this survey is organized as follows. Section 2 details the paper collection process and the inclusion criteria for selecting the articles included in this survey. Section 3 introduces the concepts and the notation we will use throughout this work. Section 4 formally defines the problem of hypergraph representation learning (see Section 4.1), describes a taxonomy of the problem setting, in terms of hypergraph embedding input and output, (see Section 4.2) and discusses the problem’s inherent challenges (see Section 4.3). Section 5 categorizes the literature based on the embedding technique, describing spectral representation learning (see Section 5.1), proximity-preserving (see Section 5.2), and (deep) neural network (see Section 5.3) methods, unraveling their pros and cons and comparing these three methodologies (see Section 5.4). Section 6 presents examples of applications enabled by the hypergraph embedding methods previously described. Section 7 identifies and discusses open research challenges and future directions in this field. Finally, Section 8 concludes this survey.

2 METHODOLOGY

This section clarifies the data collection process and the inclusion criteria for the selected articles.

2.1 Data collection

We conducted our literature review by in-depth reading, interpreting, and categorizing articles addressing the problem of generating a low-dimensional representation of a hypergraph. We carried out our search in the Scopus database since it represents a comprehensive and accurate database of peer-reviewed research articles in the fields relevant to this survey. Specifically, we submitted the following query: TITLE-ABS-KEY ("hypergraph embedding" OR "hypergraph learning" OR "hypergraph representation learning" OR "hypergraph neural network*" OR "hypergraph convolution" OR "hypergraph attention"), limiting the results to only English-written contributions. We repeated the same query by replacing the word hypergraph with the terms hyper-network and high-order, as well as trying out different spellings of these words. We did not limit either the subject areas or the publication year. The rationale behind this choice relies on two main reasons. First, hypergraphs - and, more generally, graphs - embody a tool to study emergent phenomena in a wide range of application domains. Second, hypergraphs rose to prominence only recently




Fig. 3. PRISMA chart reporting the selection process of our systematic literature review.

in the academic landscape, and the topic of hypergraph representation learning represents an even more novel research field. A total of 1338 non-duplicated articles met these criteria by June 2022.

2.2 Inclusion criteria

We followed a standard two-step selection process to pick the final set of articles to include in this survey. First, we screened the original set by filtering each article based on its title and abstract. In this phase, we removed all articles related to (higher) student education and high-order neural networks [104, 137]. In the second step, we filtered the remaining articles based on their content, removing out-of-scope articles, such as hypergraph-regularized methods. After this process, the number of articles was narrowed down to 102. Figure 3 schematically shows the selection process, while Figure 4 provides an overview of the venues where the selected articles have been published.

3 FUNDAMENTALS

This section introduces the concepts we will use throughout the paper, from formally defining hypergraphs to describing how these structures can be transformed into their graph counterparts. Table 1 lists the mathematical notation and the hypergraph-related concepts that will be explicitly referred to in the remainder of this survey.




Fig. 4. Distribution of the ranking of the venues where the selected articles have been published. Conference and journal rankings have been evaluated according to [CORE 2021](#) ($A^* > A > B > C > \text{not ranked}$) and [Scimago](#) ($Q1 > Q2 > Q3 > Q4$), respectively.

3.1 Hypergraphs

A **hypergraph** is an ordered pair $H = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of nodes¹, and \mathcal{E} is the set of hyperedges (Figure 5a). Each hyperedge is a non-empty subset of nodes. The structure of a hypergraph is usually represented by an **incidence matrix** $\mathbf{H} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{E}|}$, with each entry $\mathbf{H}(v, e)$ indicating whether the vertex v is in the hyperedge e , i.e., $\mathbf{H}(v, e) = \llbracket v \in e \rrbracket$.

The **degree of a vertex** v and the **degree of a hyperedge** e are defined as

$$\deg(v) = \sum_{e \in \mathcal{E}} \mathbf{H}(v, e) \quad \text{and} \quad \delta(e) = \sum_{v \in \mathcal{V}} \mathbf{H}(v, e),$$

respectively. $\mathbf{D}_v \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ and $\mathbf{D}_e \in \mathbb{R}^{|\mathcal{E}| \times |\mathcal{E}|}$ indicate their corresponding diagonal matrix. When all hyperedges have the same degree k , i.e., $\delta(e) = k \forall e \in \mathcal{E}$, we say that H is a **k-uniform** hypergraph. k -uniform hypergraphs also have a tensor representation [101, 205]. As for classical networks, both nodes and hyperedges may be of more than one type; in such cases, the hypergraph is **heterogeneous**.

The **dual** H^* of a hypergraph H is the hypergraph constructed by swapping the roles of nodes and hyperedges, i.e., $H^* = (\mathcal{V}^*, \mathcal{E}^*)$ where $\mathcal{V}^* \equiv \{i \mid e_i \in \mathcal{E}\}$ and $\mathcal{E}^* \equiv \{\{i \mid v \in e_i \in \mathcal{E}\} \mid v \in \mathcal{V}\}$.

In a **weighted** hypergraph, denoted by a tuple $H = (\mathcal{V}, \mathcal{E}, \mathbf{W})$, each hyperedge $e \in \mathcal{E}$ has a weight $w(e)$, representing the importance of that relation in the whole hypergraph. $\mathbf{W} \in \mathbb{R}^{|\mathcal{E}| \times |\mathcal{E}|}$ denotes the diagonal matrix of the hyperedge weights, i.e., $\text{diag}(\mathbf{W}) = [w(e_1), w(e_2), \dots, w(e_{|\mathcal{E}|})]$.

In a weighted hypergraph, the degree of a vertex v is defined as $\deg(v) = \sum_{e \in \mathcal{E}} w(e) \mathbf{H}(v, e)$. A non-weighted hypergraph can be seen as a special case of the weighted one where $\mathbf{W} = \mathbf{I}_n$.

¹In this survey, we use the terms node and vertex interchangeably.

3.2 Hypergraph to graph transformations

Hypergraphs have usually been converted into a corresponding graph representation in the literature. Even though such transformation has often been preferred over hypergraphs, especially for computation convenience and the easiness of dealing with graphs rather than higher-order structures, this process may either bring a loss of information or introduce redundant vertices/edges regarding the original hypernetwork structure. The typical transformation of a hypergraph into a graph relies on its two-section, incidence, or line graph representation [28]. Figure 5 shows a toy hypergraph and its corresponding two-section, incidence, and line graphs.

Two-section graph. The *two-section* or *clique graph* of H is the graph, denoted with $[H]_2$, whose vertices are the vertices of H and where two distinct vertices form an edge if and only if they are in the same hyperedge of H (see Figure 5b). In other words, each hyperedge of H appears as a complete sub-graph in $[H]_2$.

The major drawback of using such a transformation is that clique graphs completely lose the notion of groups since pairwise connections substitute each high-order interaction. Consequently, we have a high probability of materializing

Category	Symbol/Concept	Interpretation
General	\mathbf{v}	Bold-faced lowercase letters are used to identify vectors.
	\mathbf{A}	Bold-faced capital letters are used to identify matrices.
	$A(i, j)$	The indexed notation identifies the j -th element in the i -th row of the matrix \mathbf{A} .
	\mathcal{A}	Calligraphic capital letters are used to identify sets.
	$ \mathcal{A} $	The cardinality (i.e., number of elements) of the set \mathcal{A} .
	$\llbracket P \rrbracket$	Indicator function: it equals 1 if the predicate P is true; 0 otherwise.
	$[n]$	Set of integers from 1 to n , i.e., $\mathbb{N}_{\leq n}$.
	\mathbf{I}_n	The n by n identity matrix.
	$\text{diag}(\mathbf{A})$	Diagonal of the matrix \mathbf{A} .
	$\text{diag}(\mathbf{v})$	Diagonal matrix with its diagonal equals to \mathbf{v} .
Hypergraph-related	\mathcal{V}	Vertex set, where $ \mathcal{V} = n$.
	\mathcal{E}	Hyperedge set, with $ \mathcal{E} = m$.
	$\mathcal{E}_v \equiv \{e \in \mathcal{E} \mid v \in e\}$	Hyperedges containing the vertex v .
	\mathbf{H}	Incidence matrix of the hypergraph H .
	$\kappa(v)$	Degree of a vertex v .
	$\delta(e)$	Degree or cardinality of a hyperedge e .
	$w(e)$	Weight of a hyperedge e .
	\mathbf{W}	Diagonal matrix of hyperedge weights.
	\mathbf{D}_v	Diagonal matrix of node degrees.
	\mathbf{D}_e	Diagonal matrix of hyperedge cardinalities.
	H^*	Dual hypergraph H^* of H , constructed by swapping the role of nodes and hyperedges.
	k -uniform	Hypergraph in which each hyperedge has cardinality k .
Representation learning	\mathbf{L}	Hypergraph Laplacian.
	\mathbf{X}	Nodes' feature matrix.
	σ	Non-linear activation function.
	Θ, Φ	Learnable neural network parameters.

Table 1. Summary of the notation used throughout this survey.




Fig. 5. Hypergraph to graph transformations.

interactions that did not exist in the original hypergraph. This intuitive concept of losing the notion of groups is formalized by the fact that different hypergraphs can be transformed in the same clique graph; hence, we cannot uniquely reconstruct the original hypergraph from its clique graph. Further, clique graphs can yield computational issues as each hyperedge of size k is transformed into $\frac{k \times (k-1)}{2}$ edges.

Incidence graph. The *incidence graph* or *star expansion* of H is the bipartite graph $I(H) = (\mathcal{V}, \mathcal{E}, \mathcal{E}')$, where $v \in \mathcal{V}$ and $e \in \mathcal{E}$ are adjacent if and only if $v \in e$, i.e. $\mathcal{E}' \equiv \{(v, e) \mid v \in e\}$ (see Figure 5c).

Bipartite graph representations effectively describe group interactions (see Figure 5c). In this model, one vertex set corresponds to the hypergraph's vertices, the other to the hyperedges. Hence, a link in this graph connects a vertex to the interactions - of arbitrary order - in which it takes part. However, also bipartite graphs have a critical shortcoming inherent in their structure. Vertices in the original system do not interact directly anymore as the interaction layer always mediates their relationship. This interaction structure implies that any measure or dynamic process defined on the bipartite representation must consider this additional complexity.

Line graph. The *line graph* of H is the graph $L(H) = (\mathcal{V}', \mathcal{E}')$ such that $\mathcal{V}' \equiv \mathcal{E}$, and $\{i, j\} \in \mathcal{E}', i \neq j \iff e_i \cap e_j \neq \emptyset$ (see Figure 5d). Essentially, each hyperedge is transformed into a node and an edge is added between two nodes if the corresponding hyperedges intersect in the original hypergraph.

As happens for clique graphs, distinct hypergraphs can have identical line graphs as they lose the information about the composition of each hyperedge, storing only whether two hyperedges intersect but not in which manner. Further, sparse hypergraphs can yield relatively dense line graphs as a vertex of degree d in the hypergraph yields $\binom{d}{2}$ edges in its line graph.

4 THE HYPERGRAPH REPRESENTATION LEARNING PROBLEM

This section formally introduces the hypergraph representation learning problem to then detail the problem setting in terms of input and output. Finally, it discusses the intrinsic challenges of learning latent representations of a hypergraph.

4.1 Problem Formulation

Definition 4.1. Hypergraph embedding. For a given hypergraph $H = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of nodes and \mathcal{E} is the set of hyperedges, a hypergraph embedding is a mapping function $\Phi : \mathcal{V} \rightarrow \mathbb{R}^{|\mathcal{V}| \times d}$, where $d \ll |\mathcal{V}|$, such that Φ

defines the latent representation (a.k.a. embedding) of each node $v \in \mathcal{V}$, which captures certain network topological information in \mathcal{E} .

In other words, given a hypergraph $H = (\mathcal{V}, \mathcal{E})$ and a predefined embedding dimension d , with $d \ll |\mathcal{V}|$, the problem of hypergraph representation learning (a.k.a. hypergraph embedding) is to map H into a d -dimensional vector space (a.k.a. latent space), in which the structural properties of H are preserved as much as possible. When features are attached to nodes/hyperedges, the learned latent representation should also encode such additional information. Following this definition, each hypergraph is represented as either a d -dimensional vector (for a whole hypergraph) or a set of d -dimensional vectors, with each vector representing the embedding of part of the hypergraph such as nodes, hyperedges, or substructures.

4.2 Problem setting

In this section, we compare existing hypergraph representation learning literature from the perspective of the problem setting, describing different types of input/output and the specific characteristics of each setting.

4.2.1 Input setting. In this survey, we analyze the hypergraph embedding input along six axes: the nature of the high-order relation, its directionality and size, the temporal dimension, whether nodes have attached additional information, and whether the hypergraph is converted into a graph. Figure 6 outlines the considered input settings. Next, we introduce each category and summarize its specific characteristics.

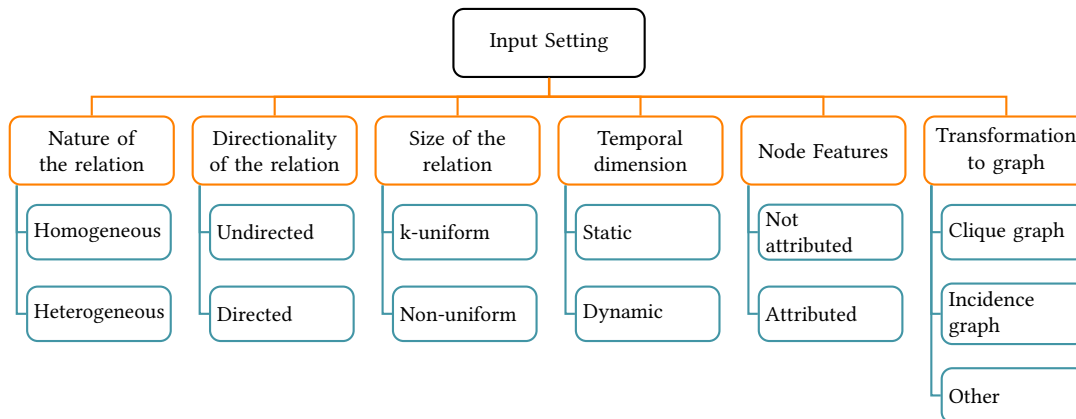


Fig. 6. Taxonomy of input settings.

Nature of the relation. As happens for graphs, hypergraphs can encode relations among nodes of one or more types. Similarly, hyperedges may represent a single or different kinds of possible interactions.

Homogeneous hypergraphs. Homogeneous hypergraphs represent the most trivial input setting, as both nodes and hyperedges belong to a single type. In this context, each hyperedge tells us that a given subset of nodes shares a common property or feature. Homogeneous hypergraphs have been widely used to model the most various user-item relations for recommendation purposes (e.g., [38, 95, 190, 206, 211]), citation networks for the link prediction task (e.g., [10, 55, 57, 97, 195]), and other relational data for classification problems (e.g., [99, 138, 150, 193, 221]). In the literature, it is also possible to find homogeneous hypergraphs built upon non-relational data [63],

where, for instance, hyperedges represent vertex connections based on some notion of distance in the feature space. In these cases, such structures are usually abstract, possibly multimodal features for object classification (e.g., [55, 118, 134, 183, 202]), traffic or passenger flow forecast [120, 177, 178], and gas or taxi demand [202].

Homogeneous hypergraphs are often employed in their weighted version to convey information about the importance of a relation through hyperedge weights (e.g., [55, 91, 118, 139, 153]). The intuition behind this choice is that the weights should drive the algorithm to be more accurate in learning the embedding of nodes within more important hyperedges.

The biggest challenge of learning a vector representation of a homogeneous hypergraph is to preserve its connectivity patterns in the latent space as only structural information is available. Typically, embedding methods defined over such structures are more general and can be re-used off-the-shelf without any particular tweak as the task they are defined for has no particular constraint.

Heterogeneous hypergraphs. Hypergraphs capture the heterogeneity of the underlying hypernetwork via nodes of different types (Figure 7c) or through nodes and hyperedges of multiple kinds (Figure 7d). In the first case, all hyperedges semantically encode the same type of interaction among diverse entities. Most of the heterogeneous hypergraphs belong to this category, and they are usually employed to model relations among users, items, and domain-specific properties or actions for ranking/recommendation (e.g., [19, 38, 175, 204, 206]), link prediction (e.g., [168, 196, 199, 200, 217]), or classification tasks (e.g., [41, 168, 225]). Hypergraphs with heterogeneous nodes and hyperedges add further expressiveness as relations can be of more than one type. In such cases, hyperedges can represent different types of events [42, 71, 72], documents, tags, and annotation relations [224], users, songs, and tags interactions [105], actions on social media [159], or even components in chemical/mechanical processes [188].

There are two major challenges when dealing with the embedding of heterogeneous hypergraphs. The first challenge relates to how effectively encode different types of nodes and relations to preserve structural and semantic properties. The second problem refers to a possible imbalance of objects of different types. Further, the more complex the heterogeneous hypergraph is, the more specific the representation learning method will be. This situation often leads to embedding algorithms that are strictly related to the application task and are tricky to generalize.

Directionality of the relation. This feature refers to whether a direct interaction exists between (groups of) nodes in the underlying hypernetwork. Hence, we can find both undirected and directed hypergraphs in the literature.

Undirected hypergraphs. By far, undirected hypergraphs represent the typical input setting in the task of hypergraph representation learning. In this case, dependencies between nodes or hyperedges are not considered.

Directed hypergraphs. Conversely to graphs, there is not a standard definition of directed hypergraphs, and the notion of direction may be either applied between hyperedges (i.e., set of nodes) [59] or between nodes within the same hyperedge [8]. Yadati et al. [196] and Gao et al. [62] follow the first interpretation, proposing an embedding algorithm working on directed hypergraphs based on the definition of Gallo et al. [59]. In this case, the head and the tail of the directed relation are embodied by two hyperedges (Figure 7e). Luo et al. [120] and La Gatta et al. [105] exploit the second interpretation, grounding their embedding technique on the definition proposed by Ausiello et al. [8]. This time the directionality of the interaction is defined within a single hyperedge, which

consists of two node sets (head and tail, Figure 7f). Liao et al. [112] propose a generalized definition of this concept, in which each directed hyperedge comprises a sequence of node sets.

Regardless of the definition of directionality used, the most critical challenge of learning a latent representation of a directed hypergraph is how to incorporate hierarchy and reachability to preserve the asymmetric transitivity in the embedding space [20, 135].

Size of the relation. This property refers to the cardinality of each hyperedge and is strictly related to the nature of the relations the hypergraph is modeling. Hyperedges can represent interactions either between a fixed or unbounded number of nodes.

k-uniform. In k -uniform hypergraphs, each hyperedge has cardinality k and usually models an existing interaction between k nodes of possible k different types (Figure 7b). Put differently, each hyperedge represents an indecomposable relation between heterogeneous entities that cease to exist when one or more components disappear. In this case, hyperedges represent relations between user mobility data and friendships (linking social, semantic, temporal, and spatial information) [167, 199, 200], or interactions between users, items, and domain-specific properties, e.g., [38, 89, 101, 168, 172]. Sometimes, a single hypergraph encodes k -way relationships of different sizes and/or types, e.g., [167, 199, 200, 206, 224].

k -uniform hypergraphs introduce more constraints regarding how the domain of interest should be modeled, but, at the same time, they ease the definition of hyperedge-related tasks, such as link prediction (see Section 4.3).

Non-uniform. In non-uniform hypergraphs, each hyperedge encodes a relation among an arbitrary number of homogeneous or heterogeneous nodes. In contrast with k -uniform hypergraphs, a hyperedge can continue to exist even when one or more nodes are removed from the network.

Heterogeneous and non-uniform hypergraphs (Figure 7c) have been used to model event data [42, 71, 72], or, as in the previous case, interactions between users, items, and domain-specific properties (e.g., [19, 52, 159, 175, 217]). Homogeneous and non-uniform hypergraphs (Figure 7a) have been exploited to represent similarities or shared properties of relational and non-relational data (e.g., [118, 142, 150, 221]), citation networks (e.g., [10, 55, 85, 182, 211]), and several user-item relations (e.g., [38, 95, 110, 206, 211]).

Thanks to their flexibility, non-uniform hypergraphs represent the most common input setting; however, conversely to k -uniform hypergraphs, they introduce further complexity to hyperedge-related tasks (see Section 4.3).

Temporal dimension. With their structure, hypergraphs can either capture a static view or the dynamicity of the underlying hypernetwork.

Static. Static hypergraphs represent the most common input setting given the early age of this research field. Static hypergraphs can either model existing connections at a fixed moment (e.g., [167, 168, 199, 200]), or node interactions over time that are aggregated into a single static snapshot (e.g., [38, 72, 89, 195, 204]).

Dynamic. Real (hyper)networks are often characterized by a dynamic behavior, meaning that both nodes and (hyper)edges can be added or removed from the system or that labels and other properties can change over time [16]. In the literature, a dynamic hypergraph is usually represented by a sequence of static hypergraphs, i.e., $H = \langle H(t_0), \dots, H(t_{T-1}) \rangle$, where $H(t_k) = (V(t_k), E(t_k))$ is a static hypergraph with timestamp t_k , with $k \in \{0, \dots, T-1\}$, T is the number of snapshots, $V(t_k)$ is the node set at timestamp t_k and $E(t_k)$ is the hyperedge set including all edges within the period $[t_k, t_{k+1})$ [110, 175].

Temporal hypergraphs have been used to model users’ preferences over time in recommendation tasks [81, 98, 110, 175], user trust relationships for rumor detection [161], or for time series forecasting [153, 177, 178, 188, 202].

Developing an approach to embedding dynamic hypergraphs implies dealing with a series of challenges that arise when the temporal component comes into play [16]. Such methods should consider (i) how to model the time domain (discrete-time or continuous-time), (ii) which dynamic behaviors have to be embedded, and (iii) which temporal granularity will be represented in the vector space.

Node features. This category refers to whether additional semantic information is associated with the nodes as input to the embedding process.

Not attributed. In this setting, nodes convey only structural information through their connectivity patterns.

Attributed. Besides the structural patterns, nodes can carry additional information about their nature in the form of feature vectors. These vectors may encode information related to user (e.g., [83, 98]) or item features (e.g., [47, 81, 190]), pre-trained word embeddings [46], image spectral features (e.g., [118, 134, 162]), and values derived from sensors (e.g., [120, 177, 178]), to name a few examples.

Although the presence of extra attributes can boost the performance of the intended task [60], this auxiliary information may not be trivial to embed, especially if it comes in a non-vector form.

Transformation to graphs. In the literature, the most common approach to deal with the high-order nature of hypergraphs is to split the single relation encoded by a hyperedge into a set of pairwise interactions. As described in Section 3.2, hypergraphs can be either transformed into clique graphs, incidence graphs, or in some intermediate representation. Considerations about the drawback of each transformation can be found in Section 3.2.

Clique graphs. Transforming a hypergraph into the corresponding clique graph means instantiating a direct interaction between each pair of nodes in a given hyperedge. Such connections can be materialized, for instance, when the hypergraph adjacency matrix is exploited to consider linking patterns (e.g., [26, 146, 150]), or used to compute the pairwise similarity of two embedding vectors of nodes within the same hyperedge (e.g., [42, 72, 167, 200]). The clique graph transformation is implicitly used by most of the hypergraph convolutional methods based on the Message Passing Framework as thoroughly described in Section 5.3.

Incidence graphs. In an incidence graph, hyperedges are represented as nodes, and the high-order relation of the nodes within a hyperedge is mediated by these special nodes via pairwise links.

From a computational perspective, incidence graphs let the information flow from nodes to hyperedges and back. For this reason, this graph transformation is implicitly used by the Message Passing techniques based on the two-stage update procedure (e.g., [7, 90, 156, 166])², in which, in a single convolutional layer, the messages pass from nodes to the hyperedges, where they are aggregated, and then back from the hyperedges to the nodes (see Section 5.3).

Line graphs. Line graphs capture the relationship between the edges of a hypergraph. This representation has been used to preserve the indecomposability of the high-order relations encoded by hyperedges in Bandyopadhyay et al. [15], which designed a hypergraph convolutional network relying on an implicit transformation to weighted line graphs. Xia et al. [190] also introduce a hypergraph transformation to a line graph to alleviate the sparsity

²Hyper-SAGE [7] does not appear in our statistics because not peer-reviewed.

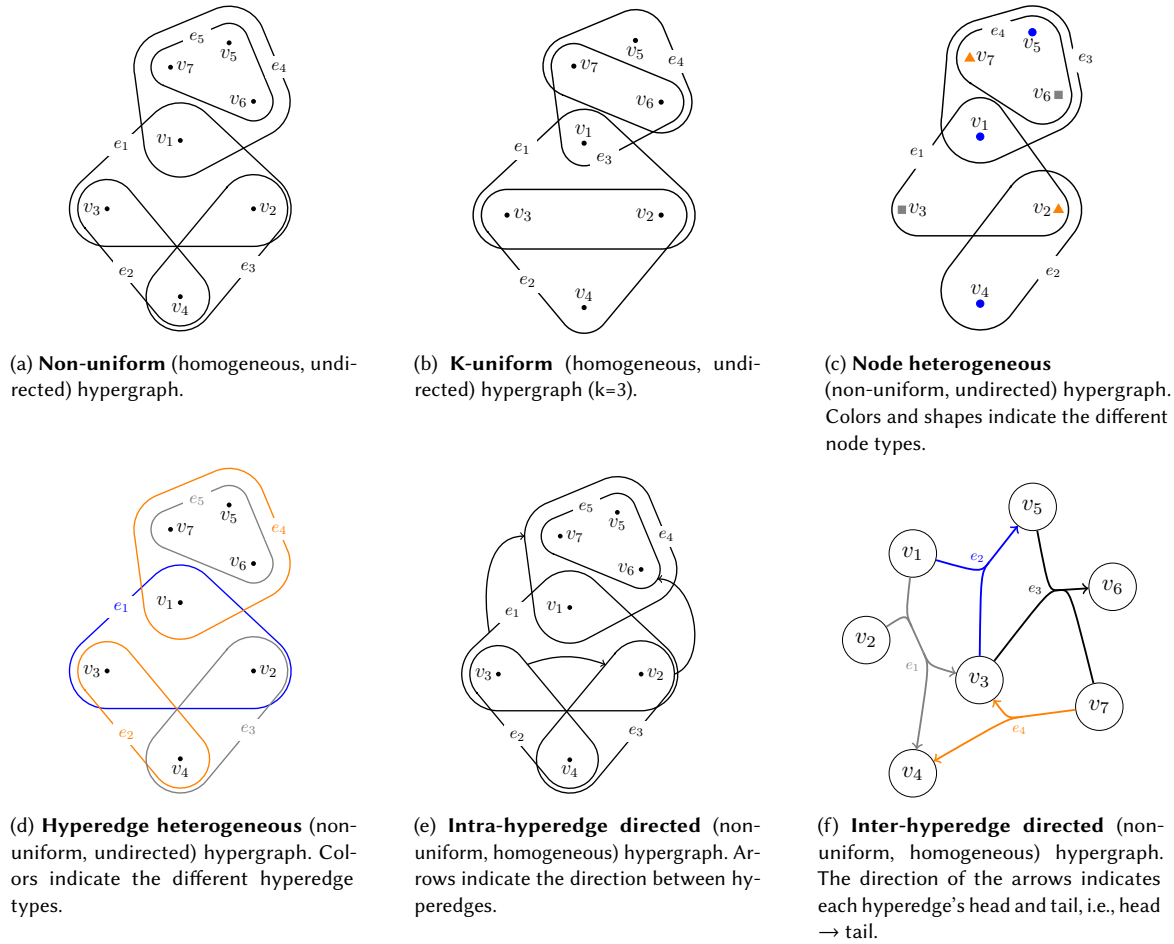


Fig. 7. Types of input hypergraphs. An example of a dynamic hypergraph is not shown because it is usually represented as a sequence of static hypergraphs.

problem in recommendation tasks. In LE [198], the line graph transformation is applied to leverage standard graph convolutional networks.

Other transformations. Other approaches convert the hypergraph into a modified form of the transformations described in Section 3.2 to better capture structural properties and use them as an implicit regularizer [195], alleviate the sensitivity towards the node distribution [141], or model domain-specific characteristics [38]. For instance, Pu and Faltings [141] construct a directed weighted line graph in which each hyperedge is replaced with two nodes (positive and negative), and an edge is added between positive and negative node pairs based on the same *ratio* of line graphs. Yadati et al. [195] propose a modified version of the classical star expansion. Also in this case, each hyperedge is replaced by two nodes, joined by an edge, and linked to all nodes originally contained in the corresponding hyperedge. A third example can be found in the work of Chen et al. [38], where

the authors define a revised version of the clique graph converting each hyperedge (of a 3-uniform hypergraph) into two pairwise interactions.

4.2.2 Output setting. The output of hypergraph embedding techniques³ is a (set of) low dimensional vector(s) representing (part of) a hypergraph. The embedding output is strictly task-driven, and finding the most suitable embedding type is crucial to meeting the specific application’s needs. In this survey, we consider *node* and *hyperedge* as embedding output settings. The most common embedding output is node embedding, while only a few approaches propose hyperedge embedding. Currently, no method deals with whole-hypergraph nor hypergraph substructures embedding (see Section 7 for more details). It is worth clarifying that there is no one-to-one correspondence between the specific embedding output setting and the addressed task since node (resp. hyperedge) embeddings can be used to evaluate hyperedge (resp. node)-related tasks.

Node embedding. Node embedding represents each node as a vector in a low-dimensional space. The underlying idea is to learn similar latent representations for nodes that are close in the original hypergraph. In practice, this concept of closeness often refers to nodes included in the same hyperedge. When nodes have additional features attached, both structural and semantic closeness are considered in the embedding process. As the most common embedding output setting, node embeddings have been used in a variety of downstream hypergraph analytic tasks, such as clustering (e.g., [42, 150, 221]), classification (e.g., [138, 159, 193]), regression (e.g., [153, 177, 202]), link prediction (e.g., [52, 85, 217]), and recommendation (e.g., [175, 176, 211]). The complexity of defining a proper similarity metric heavily depends upon the properties of the input hypergraph.

Hyperedge embedding. Hyperedge embedding methods learn low-dimensional vector representations for the hyperedges. In contrast to graphs where each edge encodes a pairwise relationship, in this context, the hyperedge embedding vector needs to capture the interaction between an arbitrary number of nodes.

The most common approach is combining the embedding vectors of the nodes within the specific hyperedge, for instance, by summing [26], averaging [72], performing more complex aggregations [42], or also by learning the aggregation function through a (deep) neural network [46, 49, 52, 62, 69, 73, 90, 138, 156, 166]. Other approaches consist in *translating* the problem of hyperedge embedding in the node domain by (i) operating on the dual hypergraph (so that the nodes of the new hypergraph represent the hyperedges of the original one) [85, 99], or (ii) on the corresponding bipartite graph (in which each hyperedge is modeled as an additional node) [164].

In this survey, we consider falling under this category only those techniques that explicitly learn hypergraph embeddings. As expected, hyperedge embedding benefits edge-related tasks, such as link prediction [52, 85, 164] or link classification [138, 156]. Nonetheless, hyperedge embedding vectors are also exploited to capture contextual information to improve node-related tasks, like clustering [42], classification [42, 49, 72, 85, 161], and recommendation [166].

4.3 Challenges of hypergraph representation learning

Obtaining an accurate representation of a hypergraph into low dimensional spaces is not a trivial task [30, 36, 210]. As hypergraphs are a generalization of graphs, the following challenges are directly inherited by the graph representation learning problem. The first challenge lies in finding the optimal embedding dimension of the representation. Using a lower dimension is more resource-efficient and may also help reduce noise in the original network; on the other hand, critical information may be lost in the process. Conversely, a higher dimension representation tends to preserve more

³Here, we are considering the embedding outputs and not the output related to the specific task (e.g., classification, regression).

information at the expense of storage requirement and computation time. Further, the proper dimension depends on the input hypergraph as well as the application domain. The second issue relates to choosing the right hypergraph property to embed, as it can be reflected by node features, link structures, or meta-data information. Again, determining which feature is the most suitable for the task is strictly domain-dependent. Such obstacles relate to preserving both the structure and the rich content that may be attached to the network elements [36, 210]. A third challenge resides in the data's nature: due to countless reasons - such as privacy or legal restrictions, among others - the problem of data sparsity may corrupt both the structure and the additional content of a network. This condition may thus lead to extra difficulties in discovering structural-level relatedness and semantic similarities between vertices not directly connected [210]. In the case of heterogeneous networks, the hypergraph representation learning task is even harder [48, 181]. In addition to the previous challenges, the core issue in this context is effectively fusing heterogeneous information to encode different types of entities and relations into latent spaces so that both structural and semantic properties are preserved. Capturing the inherent organization of such networks may require prior knowledge of the application domain to be included in the embedding process (for instance, in the design of meta-paths [69, 89]). This consideration leads to embedding techniques that are strictly application dependent and difficult to generalize and re-use in other contexts [181]. Nonetheless, embedding hypergraphs enforce to tackle two more issues deriving from the higher-order nature of these structures.

- *Capturing group relations.* Hypergraphs model many-to-many relationships among entities. Nonetheless, the existence of a hyperedge bonding of some nodes only tells us that they share a common property and does not necessarily imply direct interactions between them. For instance, a hypergraph can encode a character co-occurrence (hyper)network, in which each node is a movie character and each hyperedge is a movie scene (hence, all characters appearing in the same scene also occur in the same hyperedge) [4]. The fact that two characters belong to the same hyperedge only tells us they have appeared in the same scene, but not that they have interacted in some way. We can argue a similar observation for a heterogeneous and k -uniform hypergraph. Let us consider a hypergraph modeling a Location-based Social Network (LBSN) in which each hyperedge links a user's presence at a point of interest (POI) at a specific time along with the semantic information about the user's activity there [167, 199] (i.e., each hyperedge is a four-element set). In this case, each hyperedge represents an indecomposable relation between four types of entities, which all relate to the same real-world event. Removing a single of these entities would cause the hyperedge to disappear. A critical challenge of hypergraph representation learning lies in capturing such high-order relations in the embedding process. Mathematical tools able to abstract such relationships are fundamental in achieving this goal.
- *Task definition.* The key feature of hypergraphs is that each hyperedge embodies a relation that can potentially connect more than two nodes. This structural characteristic implies that all tasks defined over relations (i.e., formalized on the edges of a graph) - such as link prediction or network reconstruction - need to be generalized. The critical issue is that the definition of such tasks needs to be adapted to the specific application context. For instance, let us focus on the link prediction problem. Even for newbies, it is straightforward to think what is the goal of this task when instantiated on a graph, namely predicting whether a given relation between two nodes (will) exists. When dealing with hypergraphs, this problem takes on different meanings based on the nature of the underlying hypernetwork. Heterogeneous and k -uniform hypergraphs usually model fixed-sized relations among a specific group of entities with different types. In this case, we are interested in finding whether a relation (a subset of nodes) linking a fixed number of different elements exists (e.g., [89, 168, 200]). In the case

of non-uniform hypergraphs, the problem becomes harder: this time, there is no (a priori) constraint on the size of the hyperedge to predict (e.g., [52, 90, 193]). As a consequence, the computational complexity of the link prediction problem explodes as the number of different hyperedges is exponential (as it corresponds to the number of all possible subsets of a set).

Another direct consequence deriving from the structural nature of hypergraphs is that these structures can be exploited to also study problems on sets (being a hyperedge a subset of nodes). For this reason, tasks borrowed from set theory can be easily adapted in this context. A good example is the task of *hyperedge completion* [156], deriving from the set expansion task [208]. This problem involves finding one or more nodes that may fit within an existing hyperedge, thus completing it. Clearly, there is no direct equivalence of this problem for graphs.

5 HYPERGRAPH REPRESENTATION LEARNING METHODS

Hypergraph representation learning aims to embed hypergraphs into a low-dimensional space while preserving the original hypernetwork properties as much as possible.

We divide hypergraph embedding techniques into three macro categories: spectral learning methods, proximity preserving methods, and (deep) neural networks-based methods. The main difference between these three families lies in the way the representation learning problem is approached.

- The spectral learning family is the pioneer in hypergraph representation learning. Although the very first works (e.g., [26, 146]) were not focused on the learning aspects, the basic concepts and ideas apply in most of the later works (e.g., [221]). Basically, spectral learning approaches define the hypergraph representation learning problem in terms of the decomposition of the Laplacian matrix (hence, the term “spectral”) of the hypergraph (Section 5.1). The output of such decomposition (a.k.a. factorization) is the (vertex) embedding matrix, where the topological proximity of the embedding vectors is ensured by the nature of the factorization.
- Proximity-preserving methods rely on a more flexible approach to the problem: vertex proximity in the embedding space is measured in terms of a similarity function (e.g., cosine similarity), and such embeddings are learned by optimizing an objective function defined to preserve the topological proximity (i.e., close vertices should be similar) along with other factors and constraints related to the given task. The flexibility of this family of techniques lies in the modularity of the overall approach (i.e., similarity function \rightarrow objective function \rightarrow optimization) that makes it easily adaptable to very specific tasks and/or input settings.
- (Deep) Neural network-based hypergraph embedding, as the name suggests, is a family of techniques based on (deep) neural networks. As it will be clear later, nowadays, deep learning-based approaches are predominant w.r.t. the other families of hypergraph learning techniques. This success is due to the well-known strengths of Deep Neural Networks (DNNs), like automatic feature learning, scalability, and generalization, to name a few. Akin DNN for graphs, a significant role is played by hypergraph convolutions that allow the generalization of well-established neural models to work on arbitrarily structured hypergraphs. Arguably, (many) proximity-preserving methods can be implemented in terms of shallow neural networks; however, in this survey, we stick to the methodology described in the original papers, and thus we consider proximity-preserving methods not based on neural networks.

In the following, we describe each approach proposed in our taxonomy, discussing its motivation, the most representative methods following the given technique, and its pros and cons. Finally, we summarize and compare the




Fig. 8. Taxonomy of hypergraph embedding methods.

hypergraph representation learning approaches. Figure 8 outlines this taxonomy. Application tasks using the methods covered in this section are discussed in Section 6.

Figure 9 provides the temporal distribution of the papers covered in this survey divided according to their category. The overall number of works from 2022 takes into account only the articles published by June of the same year (see Section 2). Spectral learning methods were the first to appear in the literature (i.e., from the '90s). Although the works before 2008 were primarily theoretical and did not strictly focus on learning hypergraph representations as we mean nowadays, their theoretical findings were essential to developing the most recent approaches. Besides rare exceptions traditionally linked to visualization tasks, spectral learning techniques fail when facing large-scale hypergraphs. This drawback should explain their disappearance in recent years. A different consideration must be made for the proximity-preserving family. While spectral learning methods are more mathematics-driven, we can consider proximity-preserving techniques to be more machine learning-driven as they follow the usual machine learning pipeline. However, the researchers' interest in these methods started when deep neural networks were already a big deal and, as for machine learning research in general, DNNs have stolen the show. The rise of DNN techniques began right after the introduction of Graph Convolutional Networks by Kipf and Welling in 2017 [103]. Over the past few years, the




Fig. 9. Temporal distribution of the articles considered in this survey.

increase in computational power and experience working with graph convolutions has allowed a rapid escalation of DNNs methods, currently representing the *de facto* technique for hypergraph representation learning.

5.1 Spectral Representation Learning Methods

Motivation. Spectral hypergraph embedding (a.k.a. spectral representation learning) methods are historically the first representation learning methods for hypergraphs. Interestingly, as far as we know, the very first discussion about vector representation learning for nodes/hyperedges dates back to the 1980s [58], where the goal was to learn a 2D representation suitable for visualization. Despite the approach to the problem being somewhat unorthodox w.r.t. today’s standards, the work in [58] captures the essence of spectral representation learning.

Generally speaking, this methodology aims to learn a low dimensional latent representation for vertices in such a way that “similar” vertices are close to each other in the learned (Euclidean) latent space. Although, in principle, vertices’ similarity could be modeled in several ways, the vast majority of the spectral hypergraph embedding methods define it in terms of the number of common incident hyperedges. The name “spectral” stems from the usage of concepts related to spectral (hyper)graph theory [32, 44], specifically, the relationships between the eigenvectors and eigenvalues of the (hyper)graph Laplacian matrix. As we will see, the common denominator of this family of methods is the definition of the node embeddings’ features as the eigenvectors corresponding to the minimal eigenvalues of the Laplacian matrix.

Methods. Before diving into the characteristics of each spectral embedding technique, we provide a general overview of the methodology.

5.1.1 Overview. Let $H(\mathcal{V}, \mathcal{E}, \mathbf{W})$ be a weighted hypergraph, and let $\mathbf{X} \in \mathbb{R}^{n \times d}$ be the matrix where each row $\mathbf{x}_i \in \mathbb{R}^d$, with $i \in [n]$, is the latent representation of the vertex $v_i \in \mathcal{V}$ that we aim to learn. Ideally, the learned vertices’ representation should encode the structural information of the hypergraph; thus, similar vertices should be mapped onto nearby points in the latent space. If we only consider information about the hypergraph structure, the similarity between two nodes $v_i, v_j \in \mathcal{V}$ can be expressed as the (normalized) sum of the weights of the common incident hyperedges. Given this notion of similarity, we can define an optimization problem whose objective is to minimize the weighted Euclidean distance between nodes sharing the same hyperedges. Formally:

$$\arg \min_{\mathbf{X}} \frac{1}{2} \sum_{v_i \in \mathcal{V}} \sum_{v_j \in \mathcal{V}} \sum_{e \in \mathcal{E}} \mathbb{1}[v_i \in e] \mathbb{1}[v_j \in e] \frac{w(e)}{\delta(e)} \left\| \frac{\mathbf{x}_i}{\sqrt{\kappa(v_i)}} - \frac{\mathbf{x}_j}{\sqrt{\kappa(v_j)}} \right\|^2 = \underbrace{\text{tr}(\mathbf{X} (\mathbf{I}_n - \mathbf{D}_v^{-\frac{1}{2}} \mathbf{H} \mathbf{W} \mathbf{D}_e^{-1} \mathbf{H}^T \mathbf{D}_v^{-\frac{1}{2}}) \mathbf{X}^T)}_{L_{\text{Zhou}}} \quad (1)$$

subject to $\mathbf{X}^T \mathbf{X} = \mathbf{I}_d$ to avoid trivial solutions or arbitrary scaling factors, and where $\text{tr}(\cdot)$ is the trace operator, i.e., $\text{tr}(\mathbf{A}) = \sum_i \mathbf{A}(i, i)$. According to Zhou et al. [221], formulation (1) is the relaxed version of the (NP-complete) problem of finding the optimal normalized cut, whose minimum is obtained by setting as columns of \mathbf{X} (i.e., the features in the latent space) the d orthonormal eigenvectors corresponding to the d smallest nonzero eigenvalues of the hypergraph Laplacian L_{Zhou} .

5.1.2 Structural-only spectral embedding. The hypergraph Laplacian plays a key role in the definition of the loss function because it encodes the concept of “closeness” between nodes in the hypergraph. As a consequence, different Laplacians lead to different embeddings. In fact, the main difference between the methods discussed in this section lies in how the hypergraph Laplacians are defined. To ease the exposition of the hypergraph spectral embedding methods, we define a unified framework for learning the optimal d -dimensional embedding \mathbf{X} . The framework is based on the

seminal works of Bolla [26] and Zhou et al. [221], but it provides a factorized view of the hypergraph Laplacian, allowing instantiating the different techniques by setting those factors. Formally, the framework is defined as

$$\arg \min_{\mathbf{X}: \mathbf{X}^T \mathbf{X} = \mathbf{I}_d} \underbrace{\text{tr}(\mathbf{X}^T \mathbf{Z} (\mathbf{D}_v - \mathbf{S}_v) \mathbf{Z}^T \mathbf{X})}_{\mathbf{L}} = \text{tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}), \quad (2)$$

where $\mathbf{L} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is the hypergraph Laplacian, $\mathbf{Z} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is a normalization matrix, and $\mathbf{S}_v \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ can be interpreted as a similarity matrix between nodes. When not specified differently, we assume $\mathbf{S}_v = \mathbf{H} \mathbf{S}_e \mathbf{H}^T$, where $\mathbf{S}_e \in \mathbb{R}^{|\mathcal{E}| \times |\mathcal{E}|}$ is a function of the weight and/or cardinality of the hyperedges. So, the factor $(\mathbf{D}_v - \mathbf{S}_v)$ represents a generalization of the not normalized graph Laplacian (i.e., the graph Laplacian can be obtained with $\mathbf{S}_v = \mathbf{H} \mathbf{H}^T$). The matrix \mathbf{Z} acts as a Laplacian normalizer, and when it is not the identity (i.e., no normalization), it is equal to $\mathbf{D}_v^{-1/2}$. Normalized Laplacians are usually preferred because the influence of hub nodes (i.e., nodes contained in many hyperedges) is reduced, and their definition is tied to the probability transition matrix.

By definition, \mathbf{L} is positive semi-definite, and, as previously described, we know that the solution to (2) can be computed using the eigendecomposition of \mathbf{L} from standard results in linear algebra [21].

The problem (2) was first introduced by Bolla in [26], where it is instantiated to find the “minimal variance placement” of the vertices in the latent space, leading to a formulation with an unnormalized Laplacian, i.e., $\mathbf{Z} = \mathbf{I}_{|\mathcal{V}|}$ and $\mathbf{S}_e = \mathbf{D}_e^{-1}$. Bolla also showed that, given the optimal \mathbf{X}^* , the optimal hyperedge embeddings can be computed as $\mathbf{X}^* \mathbf{H} \mathbf{D}_e^{-1}$.

Zhou et al. [221] made a step forward by providing a general framework for spectral hypergraph partitioning. The proposed normalized hypergraph Laplacian (i.e., \mathbf{L}_{Zhou}) generalizes [26] by including hyperedge weights, i.e., $\mathbf{S}_e = \mathbf{W} \mathbf{D}_e^{-1}$, as well as node normalization, i.e., $\mathbf{Z} = \mathbf{D}_v^{-1/2}$. Closely related to Zhou’s Laplacian are the ones proposed by Zhu et al. [224] (HHE [224]) and Pu et al. [141]. In HHE [224], the Laplacian corresponds to Zhou’s Laplacian without the node normalization ($\mathbf{Z} = \mathbf{I}_{|\mathcal{V}|}$), while in Pu et al. the hyperedge degree matrix is omitted in \mathbf{S}_e . Pu et al. also present a novel Laplacian matrix (that does not fit the framework in (2)) based on the hyperedge expansion so that the learning result is invariant to the distribution of vertices among hyperedges. This novel formulation seems to get clusters of nodes with larger margins which could be beneficial in the case of clustering tasks.

In [142], Ren et al. state that Zhou’s Laplacian is not ideal for vision problems. Specifically, they argue that \mathbf{S}_e contains redundant information while the normalization is empirically ineffective. For this reason, they proposed a Laplacian analogous to the standard graph Laplacian, i.e., $\mathbf{S}_e = \frac{1}{2} \mathbf{I}_{|\mathcal{E}|}$ and $\mathbf{Z} = \sqrt{2} \mathbf{I}_{|\mathcal{V}|}$. Saito et al. [150] proposed a Laplacian strongly related to Zhou’s where the node similarity matrix has the diagonal zeroed out, i.e., $\mathbf{S}_v \leftarrow \mathbf{S}_v - \text{diag}(\mathbf{S}_v)$. From a random walk standpoint, Saito’s Laplacian is consistent with standard graph Laplacians, whereas Zhou’s setting can be regarded as a lazy random walk with self-loops.

The direct connection with graphs is also present in Rodriguez [146], where the hypergraph Laplacian of $H(\mathcal{V}, \mathcal{E})$ corresponds to the Laplacian matrix of the weighted graph on \mathcal{V} , in which two vertices v, u are adjacent if they are adjacent in H , and the edge-weight is the number of edges in H containing both u and v , i.e., $|\mathcal{E}_u \cap \mathcal{E}_v|$. According to Rodriguez, this Laplacian has nice spectral properties; however, the impact in terms of node embeddings is unclear.

5.1.3 Spectral embedding for nodes with additional features. Despite its generality, the framework (2) assumes that nodes only carry structural information. However, there are many applications in which additional features can be attached to nodes. We can generalize the framework to include such cases and consider an initial node representation. Let $\mathbf{V} \in \mathbb{R}^{n \times f}$ be the matrix containing each node (on the rows) and its features (on the columns); then, we can rewrite

(2) as

$$\arg \min_{\mathbf{X}: \mathbf{X}^T \mathbf{V}^T \mathbf{V} \mathbf{X} = \mathbf{I}_d} \text{tr}(\mathbf{X}^T \mathbf{V}^T \mathbf{Z} (\mathbf{D}_v - \mathbf{S}_v) \mathbf{Z}^T \mathbf{V} \mathbf{X}) = \text{tr}(\mathbf{X}^T \mathbf{V}^T \mathbf{L} \mathbf{V} \mathbf{X}), \quad (3)$$

where $\mathbf{X} \in \mathbb{R}^{f \times d}$ is the projection matrix that we want to learn. So, differently from (2), here we learn a linear transformation rather than the embedding which is obtained as $\mathbf{V} \mathbf{X} \in \mathbb{R}^{n \times d}$. It is worth noticing that Equation (3) generalizes Equation (2) and they are equivalent if \mathbf{V} is the identity matrix with $f = n$.

Both the approaches proposed by Sun et al. in [162] and [158] fall under the formulation (3), while the approaches presented by Luo et al. in [118] and [119] are closely related to (3) but the optimized loss is different. Specifically, Lou et al. use Zhou’s Laplacian, but embeddings are learned in a supervised fashion where the supervision is used to define an objective function that simultaneously ensures high intra-class compactness and high inter-class separability. Very similar approaches to Luo’s are proposed by Yuan et al. [207], and Huang et al. [91].

Remarks. As demonstrated by Agarwal et al. [1], all hypergraph Laplacians introduced in this survey do not consider the higher-order nature of the interactions encoded by hypergraphs as, in practice, those are manipulated as graphs. Precisely, with the appropriate weighting function, [221, 224]’s Laplacians correspond to the Laplacian of the hypergraph star expansion, while [26, 146, 150]’s Laplacians are equivalent to the Laplacian of the hypergraph clique expansion. Agarwal et al. further showed that, in a linear setting, these two transformations are equivalent.

These embedding methods benefit from spectral theory on hypergraphs, making their definition elegant and easy to implement since they are based on standard linear algebra operations. However, the overall spectral framework is not flexible enough to easily handle the injection of node and hyperedge heterogeneity, and it is not designed to deal with temporal dependencies. Such settings can be tackled in a naive way by constructing different hypergraph laplacians based on the node/hyperedge types [224] or by building discrete hypergraph snapshots and learning node embeddings by solving the generalized eigenproblem (as for dynamic graphs [192]), respectively. Further, as in the case of graphs, spectral theory on hypergraphs does not straightforwardly translate to directed hypergraphs (whichever their definition is, see Section 4.2.1), and solutions to make the adjacency matrix symmetrical would add further (computational)

Method	Input HG	Laplacian = $\mathbf{Z} (\mathbf{D}_v - \mathbf{S}_v) \mathbf{Z}^T$	Learning Task	Venue	Year
Bolla [26]	Hom, NU	$\mathbf{S}_e = \mathbf{D}_e^{-1}$ and $\mathbf{Z} = \mathbf{I}_{ \mathcal{V} }$	-	Discrete Mathematics	1993
Rodriguez [146]	Hom, NU	$\mathbf{D}_r = \text{diag}(\mathbf{H} \mathbf{W} \mathbf{H}^T - \mathbf{D}_v)$, $\mathbf{Z} = \mathbf{D}_r^{-1/2}$, $\mathbf{S}_e = \mathbf{W}$, $\mathbf{S}_v = \mathbf{H} \mathbf{S}_e \mathbf{H}^T - \mathbf{D}_r$	-	Linear & Multilinear Algebra	2002
Zhou et al. [221]	Hom, NU	$\mathbf{S}_e = \mathbf{W} \mathbf{D}_e^{-1}$ and $\mathbf{Z} = \mathbf{D}_v^{-1/2}$	Clustering	NeurIPS	2007
Ren et al. [142]	Hom, NU	$\mathbf{S}_e = \frac{1}{2} \mathbf{I}_{ \mathcal{E} }$ and $\mathbf{Z} = \sqrt{2} \mathbf{I}_{ \mathcal{V} }$	Clustering	SSPR	2008
Sun et al. [158]	Hom, NU, A	Zhou’s Laplacian	Classification	KDD	2008
Pu et al. [141]	Hom, NU	$\mathbf{S}_e = \mathbf{W}$ and $\mathbf{Z} = \mathbf{D}_v^{-1/2}$ (*)	Classification	ECML PKDD	2012
Yuan et al. [207]	Hom, NU, A	Zhou’s Laplacian (*)	Classification	IEEE GRSL	2015
Zhu et al. [224]	Het, NU	$\mathbf{Z} = \mathbf{I}_{ \mathcal{V} }$ and $\mathbf{S}_e = \mathbf{W} \mathbf{D}_e^{-1}$	Recommendation	Neurocomputing	2016
Huang et al. [91]	Hom, NU, A	Zhou’s Laplacian (*)	Dimensionality reduction	Neurocomputing	2016
Sun et al. [162]	Hom, NU, A	$\mathbf{Z} = \mathbf{I}_{ \mathcal{V} }$ and $\mathbf{S}_e = \mathbf{W} \mathbf{D}_e^{-1}$	Classification	Remote Sensing	2017
Saito et al. [150]	Hom, NU	$\mathbf{Z} = \mathbf{D}_v^{-1/2}$, $\mathbf{S}_e = \mathbf{D}_e^{-1} \mathbf{W}$, $\mathbf{S}_v = \mathbf{A} - \text{diag}(\mathbf{A})$, $\mathbf{A} = \mathbf{H} \mathbf{S}_e \mathbf{H}^T$	Clustering	AAAI	2018
Luo et al. [118]	Hom, NU, A	Zhou’s Laplacian (*)	Classification	IEEE Trans. on Cybernetics	2019
Luo et al. [119]	Hom, NU, A	Zhou’s Laplacian (*)	Classification	JSTARS	2020

Table 2. *Spectral-based methods.* As tasks, we only consider those directly using the node embeddings (standard output among these methods). All methods are transductive and work in a static and undirected input setting. Hom/Het stand for homogeneous/heterogeneous; NU stands for non-uniform; A stands for attributed nodes. (*) The formulation is related but does not directly fit the framework (3).

complexity to an already demanding problem [219]. In addition, spectral methods scale poorly with $|V|$ as they require the in-memory storage of all the matrices involved in (1), even though sparse representations (of sparse hypergraphs) may help alleviate this problem. Still, this requirement hugely hinders their applicability to very large hypergraphs. This scaling issue explains why most spectral embedding methods were proposed in early 2000, while nowadays, more scalable approaches are preferred, e.g., neural network-based methods.

Table 2 provides an overview of the spectral methods described in this section. We provide the implementation of most of these methods at <https://github.com/alessant/HEE>.

5.2 Proximity-Preserving Methods

Motivation. Proximity-preserving (PP) algorithms aim to design a proper similarity function able to catch various structural and semantic information conveyed by the hypergraph. The similarity can be naturally computed in graphs based on first-order or second-order proximity information. While the former tells how similar two nodes are based on the weight of the link between them, the latter compares the similarity of the nodes' neighborhood structures [30]. When it comes to hypergraphs, a desirable property of the embedding is to simultaneously retain the proximity among the nodes in the same group, as a hyperedge encodes high-order relationships that are not necessarily meaningful when fractured into pairwise links [204].

Methods. Proximity-preserving and spectral embedding algorithms share the idea of retaining the nodes' topological closeness; however, the approach to the problem is very different. PP techniques are less theoretically grounded, but the methodology is much more flexible and easy to adapt to the specific task at hand.

Although it is impossible to provide a single sensible framework general enough to cover all the methods falling in this category, we can highlight the concepts and ideas shared among most, if not all, of them.

Similarity as a function of the dot-product. In the case of Euclidean latent spaces, the proximity/closeness of two vectors is commonly measured using the dot-product [30]. Based on this general idea, PP methods designed for hypergraphs generally aim to jointly optimize the tuple-wise (or n -wise) proximity of the nodes in a hyperedge via the dot-product (HEBE(-PO/PE) [71, 72], FOBE and HOBE [165]), or one of its variants, like the cosine similarity (Lbsn2Vec [199], Lbsn2Vec++ [200], and MSC-LBSN [167]).

Negative sampling. Regardless of the specific proximity function used, optimizing the similarity of related nodes comes with the risk of overfitting. In order to avoid overfitting, PP methods need (at least) to consider negative hyperedges (i.e., hyperedges in $2^E \setminus \mathcal{E}$) by maximizing the distance of the incident nodes. Given the combinatorial explosion of the number of non-existing hyperedges, negative sampling represents the most common technique to solve this problem (Lbsn2Vec(++) [199], HGE [204], HEBE(-PE/PO) [71, 72], FOBE, and HOBE [165]).

Objective function. The objective function presents a similar structure across all PP methods. Specifically, we can identify three primary sub-objectives:

- (1) maximization of the node embeddings' similarity for nodes that are topologically or semantically close to each other;
- (2) minimization of the node embeddings' similarity for the negative samples (when present);
- (3) maximization of a task-specific sub-objective that can be intrinsically embedded into the objective (1). An example of such a sub-objective is the maximization of the *inter-event proximity* used in Event2Vec [42], which describes the pairwise proximity between hyperedges.

No regularization. Surprisingly, most of the PP methods do not employ any kind of regularization over the learned embeddings (with the only exceptions of MSC-LBSN [167] and HOBE [165]). For methods based on the cosine similarity, e.g., Lbsn2Vec(++) [199], the regularization may not be necessary since the magnitude of the embedding vectors does not impact the cosine of the angle between them. However, it is unclear whether the lack of regularization may lead to overfitting for methods based on the dot-product, like HEBE [71, 72], Event2Vec [42], FOBE [165], and HGE [204].

Despite being similar at a higher level, each PP technique is characterized by the optimization objective, usually optimized via Stochastic Gradient Descent, which shapes how the node/hyperedge embeddings are learned. The optimization criterion can take different forms like similarity maximization (e.g., Lbsn2Vec(++) [199], HGE [204], MSC-LBSN [167]), mean squared error (e.g., HOBE [165]), Kullback-Leibler divergence (e.g., FOBE [165], Event2Vec [42]), or Bayesian Personalized Ranking [144] (e.g., HEBE(-PO/PE) [71, 72]). Clearly, the loss function takes into account the characteristics of the task to solve.

PP methods also differ in how they transfer the high-order relation conveyed by a hyperedge and whether they explicitly learn a representation for the hyperedge itself. Almost all methods of this family maximize the embedding similarity of nodes within the same hyperedge by pairwise computing the cosine similarity (Lbsn2Vec(++) [199], MSC-LBSN [167]) or the dot-product (HEBE [71], Event2Vec [42], HOBE/FOBE [165]) between two embedding vectors. An exception is made by HGE [204], which treats a hyperedge as a single set of nodes. Specifically, HGE [204] generalizes the standard dot-product operation between two vectors to an indefinite number of vectors for computing the similarity score of the nodes in a hyperedge. However, this dot-product generalization, called multilinear map, is meaningful only from an algebraic perspective while having no geometric interpretation. Among all PP methods, only HEBE-PE [71, 72] and Event2Vec [42] also learn a latent representation for hyperedges. In HEBE-PE [71, 72], the hyperedge embedding is learned by maximizing the similarity between the hyperedge vector and the average representation of the nodes within that hyperedge. In Event2Vec [42], the latent hyperedge representation is given by combining via a weighted dot-product the embedding of the nodes within it.

Another important observation is that all PP methods, but Event2Vec [42], only consider the first-order proximity, which could be sub-optimal. In particular, Event2Vec [42] captures the second-order proximity by modeling hyperedge proximity. Inspired by Factorization Machines [143], the authors proposed a new operation combining all the interactions between pairs of embedding vectors in the same hyperedge. Then, hyperedge proximity is modeled with the dot-product between the embedding and the just described combination.

Remarks. Proximity-preserving methods share the same overall design: the definition of (i) a similarity measure as a function of the embeddings and (ii) a criterion to optimize for preserving the node proximity and possibly task-specific goals. This simple design makes PP methods versatile enough to adapt to peculiar tasks easily by means of specifically crafted optimization (sub-)objectives. This elementary structure also allows to naturally handle heterogeneous, non-uniform, and directed hypergraphs. In particular, the notion of directionality can be encoded within the objective function. For instance, in the case of inter-hyperedge directed hypergraphs (see Figure 7f), the criterion to optimize could enforce the embedding of the tail node(s) to have some particular characteristics (e.g., being the centroid of the embedding of all nodes within a hyperedge). The injection of dynamic information into the latent representation via a PP method deserves separate considerations. In this case, the dynamic setting would make the overall method's structure harder to design, consequently conflicting with the underlying spirit of these techniques, which tend to be straightforward and efficient. Currently, no PP method is defined over directed or dynamic hypergraphs.

Method	Input HG	Proximity Model	Learning Task	Venue	Year	Code
HEBE [71]	Het, NU	Softmax of the dot-product	Classification	ICDM	2016	https://bitbucket.org/hgui/hebe/downloads/
HEBE-PO [72]	Het, NU	As HEBE	Classification	TKDE	2017	https://bitbucket.org/hgui/hebe/downloads/
HEBE-PE [72]	Het, NU	As HEBE	Link prediction	TKDE	2017	https://bitbucket.org/hgui/hebe/downloads/
HGE [204]	Het, NU	Multilinear map	Recommendation	CIKM	2018	https://github.com/chia-an/HGE
Event2vec [42]	Het, NU	Sigmoid of the weighted dot-product	Clustering, Class.	ICDMW	2018	-
FOBE [164]	Hom, NU	Sigmoid of the dot-product	Link prediction, Recom.	MLG	2019	https://github.com/JSybrandt/HypergraphEmbedding
HOBE [164]	Hom, NU	ReLU of the dot-product	Link prediction, Recom.	MLG	2019	https://github.com/JSybrandt/HypergraphEmbedding
LBSN2Vec [199]	Het, U	Cosine similarity	Link prediction, Recom.	WWW	2019	https://github.com/eXascaleInfolab/LBSN2Vec
LBSN2Vec++ [200]	Het, U	Cosine similarity	Link prediction, Recom.	TKDE	2020	https://github.com/eXascaleInfolab/LBSN2Vec
MSC-LBSN [167]	Het, U	Cosine similarity	Link prediction, Recom.	TKDE	2022	-

Table 3. *Proximity-preserving methods*. Node embedding is the standard output. All methods are transductive and work in a static and undirected input setting. Hom/Het stand for homogeneous/heterogeneous; (N)U stands for (non-)uniform. Task: Class stands for classification, and Recom for recommendation.

In contrast with deep representation learning (Section 5.3), we can define these methods as shallow because they optimize a unique embedding vector (i.e., a single layer of abstraction) for each node/hyperedge. These approaches are inherently transductive, which prevents them from being used in inductive applications. A drawback of most of these methods is their lack of capturing structural hypergraph information besides single-hop neighbors. Moreover, since similarity is usually defined as a function of the dot-product, these approaches may fail to model the potential non-linear relationship between nodes/hyperedges. Finally, PP methods on hypergraphs often treat the hypernetwork as its clique expansion counterpart, which may lose the benefit of modeling the problem in terms of high-order interactions in the first place. Table 3 summarizes the methods described in this section.

5.3 (Deep) Neural Network Models

Motivation. Deep Learning (DL) has gained remarkable impetus in speech, language, and visual detection systems [64]. This success drove the research community to either directly apply DL models from other fields on graphs or design novel neural network models for specifically embedding graph-shaped data [30]. Hypergraph neural networks (HNNs) stem from graph NNs (GNNs), which currently embody the state-of-the-art in graph representation learning [184]. For this reason, the motivation and intuition behind GNNs also apply to hypergraph-like input data. (Deep) HNNs have been utilized to capture the high-order non-linear relationships between the nodes of a hypergraph without requiring hand-crafted features (i.e., providing end-to-end solutions). Furthermore, state-of-the-art NN frameworks (e.g., Pytorch, Keras, TensorFlow) leverage hardware acceleration (e.g., GPUs) to speed up the computation making it possible to train very deep, thus powerful, models efficiently.

Methods. Unlike (most of the) non-neural embedding techniques, the (deep) HNN framework requires having node features $\mathbf{x}_v \in \mathbb{R}^d, \forall v \in \mathcal{V}$ (usually arranged into a feature matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$) as input to the model. When node features are not available, these can be initialized using a (parametric) encoding (e.g., [196, 217]), one-hot encoding (e.g., [46, 81, 99]), or *ad hoc* initializations (e.g., [39, 69, 125, 196]). Such initial representation is then fed to a (deep) neural network that learns the node (hyperedge/hypergraph) representation either in an end-to-end fashion or in a two-step procedure. In the former case, the learning process is guided by the specific task; thus, the embeddings are optimized (according to the loss function) to solve the problem at hand better. In the latter case, the embeddings are learned in an

unsupervised fashion or by solving a fake task (e.g., [41, 85, 89]). Then, such embeddings are fed to another model, e.g., a classifier that may not be a neural network, to solve the task (e.g., [138, 217]).

In the following, we focus our description on the parts of the NN that explicitly or implicitly learn the latent hypergraph representation, giving less importance to the readout layers.

5.3.1 Hypergraph convolutional networks. Graph convolution networks are by far the most popular technique for learning graph embedding. The convolution operator on (hyper)graphs generalizes the operation of convolution from grid data to graph data [184] and can be elegantly motivated using the theory of (hyper)graph signal processing as a generalization of Euclidean convolutions to non-Euclidean data.

In the literature, (hyper)graph convolutions are usually classified between spatial and spectral based on how the convolution operator is defined [184, 218]. On the one hand, spectral graph convolution uses the Fourier transform to transform the graph signal to the spectral domain, where it carries out the convolution operation. On the other hand, spatial graph convolution aggregates the node features from the perspective of the spatial domain [213]. However, as thoroughly shown in [13, 14], this distinction is becoming less and less clear since many graph convolutions can be defined in both spectral and spatial domains. Based on this consideration, we prefer to avoid such a blurry categorization in our survey in favor of a more general framework encompassing almost all convolutional methods while still giving some space to peculiar spectral convolutions.

As Hamilton et al. [75] show, graph convolutions are a particular case of the more general and reader-friendly Message-Passing framework (MPF), whose underlying intuition is straightforward. Given an initial hypergraph representation (e.g., nodes' feature matrix $\mathbf{X}^{(0)}$), the MPF iteratively updates it according to the following process. At each iteration, every node aggregates information from its local neighborhood, and, as the iterations progress, each node embedding contains more and more information from further reaches of the hypergraph. Hence, node embeddings encode two-fold knowledge: structural and feature-based information, both deriving from iteratively gathering neighbors' representations according to the hypergraph structure. Each message-passing step is represented as a layer of a (deep) hypergraph convolutional NN (convHNN), which may act as input to the next layer.

Conversely to graphs where every edge simply connects two nodes, hyperedges in hypergraphs encode many-to-many interactions among a set of two or more nodes. Thus, to properly harness the high-order nature of hypergraphs, the propagation of the information over nodes should consider such richer relations. This intuition can be accomplished using a two-stage (a.k.a. two-level or two-step) update/aggregation process [7]:

$$\mathbf{x}_e^{(l+1)} = f_{\mathcal{V} \rightarrow \mathcal{E}} \left(\mathbf{x}_e^{(l)}, \{\mathbf{x}_v^{(l)}\}_{v \in e}; \Theta^{(l)} \right), \quad \mathbf{x}_v^{(l+1)} = f_{\mathcal{E} \rightarrow \mathcal{V}} \left(\mathbf{x}_v^{(l)}, \{\mathbf{x}_e^{(l+1)}\}_{v \in e}; \Phi^{(l)} \right) \quad (4)$$

where $\mathbf{x}_e^{(l)}, \mathbf{x}_v^{(l)}$ are the latent representations of hyperedges and nodes at layer l , respectively; $f_{\mathcal{V} \rightarrow \mathcal{E}}, f_{\mathcal{E} \rightarrow \mathcal{V}}$ are two (non-linear) parametric permutation-invariant aggregation functions known as *node-to-hyperedge* and *hyperedge-to-node* aggregation, respectively; $\Theta^{(l)} \in \mathbb{R}^{d_l \times d_{l+1}}$ and $\Phi^{(l)} \in \mathbb{R}^{d'_l \times d'_{l+1}}$ are the learnable convolutional parameters. As for standard NNs, parameters are learned via the optimization of objective functions specifically designed for the given task. The optimization is performed using backpropagation and gradient descent-based algorithms like Stochastic Gradient Descent (SGD) and Adam [102]. For simplicity, we omit the input argument H and the potential set of hyper-parameters from (4). Formulation (4) is very similar to the one proposed by Chien et al. [40] (AllSet [40]), and it generalizes many existing works (UNIGNN [90], HyperSaR [166], HyperGAT [46], HNHN [49], HGC-RNN [202], [7]). Peculiar two-stage MP processes are the ones proposed in [69], and [73]. In [69], Guan et al. define $f_{\mathcal{V} \rightarrow \mathcal{E}}$ as a *hyperedge shrinking* function in which they select the most informative (like an attention mechanism) nodes to update the hyperedge representation. In

[73], instead, one of the stages is fixed. Specifically, node embeddings are aggregated using a mean and are not updated in the MP phase. Node embeddings are learned beforehand using a Graph Convolutional NN, and their weighted mean is used as the initial hyperedge representations.

A slightly different two-stage MPF is proposed by Srinivasan et al. [156], where the nodes and the hyperedges are updated simultaneously (in (4), the updates happen sequentially) and where the aggregation functions also consider the embeddings of the second order neighbors.

While only the few works cited above rely on the two-stage MPF, most of the hypergraph convolutional NN usually resort to the hypergraph’s (weighted) clique-expansion; thus, reducing (4) to a single-stage MPF (as for graphs) first introduced by Feng et al. [55] (HGNN). In this case, the t -th (single-stage) hypergraph convolutional layer can be elegantly defined as

$$\mathbf{X}^{(l+1)} = \sigma(\overbrace{\mathbf{R}\mathbf{X}^{(l)}}^{\text{update}} \overbrace{\Theta^{(l)}}^{\text{aggregate}}), \quad (5)$$

where $\mathbf{R} \in \mathbb{R}^{n \times n}$ is the so-called *reference operator* [66], usually instantiated as a “surrogate” of the adjacency matrix, $\Theta^{(l)} \in \mathbb{R}^{d_l \times d_{l+1}}$ are the learnable convolutional parameters (with d_l the dimension of the latent representation at layer l), and σ is a non-linear activation function (e.g., sigmoid). It is worth noting that although the MPF is flexible enough to allow to send different messages across the neighbors, the convolutional layer in (5) assumes that every node sends the same message to each of its neighbors. Moreover, since the update of the node embeddings happens directly, hyperedge embeddings are not explicitly learned in the process. However, one can learn hyperedge embeddings by resorting to the dual transformation and applying the MPF (EHGNN [99]). To reduce the two-section graph connectivity and speed up the learning process, Yadati et al. [195] ((Fast)HyperGCN) proposed replacing each hyperedge with an incomplete clique via so-called mediators.

Eq. (5) covers most of the hypergraph convolutional models (discussed later), however, there are exceptions that may only fit in (4). For instance, H^2 SeqRec [110] defines an *ad hoc* convolutional operation in the hyperbolic space to alleviate the sparsity issue in hypergraphs for recommendation tasks, while DHGNN [97] and KHNN [114] use a 1D convolution and a parametric aggregation function to better model discriminative information among nodes. In [194], Yadati et al. proposed G-MPNN [194], a generalized MPF able to handle multi-relational ordered hypergraph, and MPNN-R to handle recursive hypergraphs.

There are also methods that use a non-parametric aggregation function just to propagate information in the hypergraph and to initialize node representations, such as (D)HCN [190] and MHCN [206]).

In the following, we discuss five convolutional design factors (i.e., reference operator, skip-connections, attention mechanism, gated update and spectral convolution) that impact how the information is propagated and aggregated during the learning process. These design choices may be coupled together in the same network architecture.

Reference operator. In Eq. (5), the matrix \mathbf{R} describes how embeddings are aggregated. In its simplest form, \mathbf{R} is exactly the adjacency matrix of the hypergraph clique expansion, defined as $\mathbf{H}\mathbf{H}^T - \mathbf{D}_v$ (NHP [196], SHCN [38], HHNE [19], DHCN [190] HGNN⁺ [62]). In this case, the aggregation function is the sum of the representations of the neighbor nodes.

However, a not normalized adjacency matrix may hinder the optimization process as node features may be in very different ranges and, further, feature values may increase indefinitely on deep networks. For these reasons, the adjacency matrix is commonly normalized. The row-normalized adjacency matrix becomes $\mathbf{R}_{\text{row}} = \mathbf{D}_v^{-1}\mathbf{H}\mathbf{D}_e^{-1}\mathbf{W}\mathbf{H}^T$ (HCHA [10], HHGR [211], GC-HGNN [139] and IHGNN [39] without the hyperedge-related matrices), while it has the form

$\mathbf{R}_{\text{sym}} = \mathbf{D}_v^{-\frac{1}{2}} \mathbf{H} \mathbf{D}_e^{-1} \mathbf{W} \mathbf{H}^T \mathbf{D}_v^{-\frac{1}{2}}$ if symmetrically normalized (HGNN [55], GC-HGNN [139], HGNN [47], HyperCTR [81], DualHGNN [182], DualHGNN [193], MHGNN [9], STHAN-SR [153], AdaHGNN [183], MultiHGNN [87], STHGNN [178], F²HNN [125], S²HGN [124], HGDD [136], HybridHGNN [88], HyperINF [98], MT-HGNN [177], HGNN⁺ [62], DH-HGNN [78]).

This latter formulation is usually preferred since row normalization does not consider the connectivity of the neighbor nodes and, as such, gives the same relative importance to highly and sparsely connected nodes. It is worth noticing that both normalized versions also consider self-loops (i.e., \mathbf{R} has a nonzero diagonal); thus, the representation of a given node at layer l is transformed and included in its representation at layer $l + 1$ (i.e., skip-connection). Node level normalization is also possible, and in such a case, \mathbf{W} (DHCF [95], HAIN [15]) and/or \mathbf{D}_e^{-1} (HHNE [19], MGCN [37], HyperRec [175], [188]) can be removed from the computation of $\mathbf{R}_{\text{sym}/\text{row}}$. The set of learnable parameters is shared among all nodes.

An interesting approach is devised by Liu et al. in HGCELM [115]. Specifically, the authors defined a single-layered random hypergraph convolution based on (5) with \mathbf{R}_{sym} , where the parameters are initialized randomly and used to perform a random projection (inspired by extreme learning machines [84]). Such representations are then fed to a hypergraph convolutional regression layer to predict node labels.

In both HCCF [186] and SHT [187], the hypergraph topology is not fixed but jointly learned with the embeddings. To do so, the reference operator is parametric, and it also includes a nonlinear activation function (i.e., LeakyReLU [191]). *Skip-connections.* Over-smoothing is a well-known problem of deep (hyper)graph convolutional networks [34]. Over-smoothing occurs when the information aggregated from the neighbors during the message passing starts to dominate the updated node representations. This phenomenon happens as the receptive field of the convolution grows exponentially with respect to the model depth. A possible solution to alleviate this issue is to use skip-connections (a.k.a., residual connections/layer), which try to directly preserve the node information from the previous round of message passing in the update. This concept can also be generalized by including all (or some of the) previous node representations. Hypergraph convolution can be integrated with generalized skip-connections. In this case, Eq. (5) becomes:

$$\mathbf{X}^{(l+1)} = \sigma \left(\mathbf{R} \mathbf{X}^{(l)} \Theta^{(l)} + \underbrace{\sum_{j=0}^l \mathbf{X}^{(j)} \Phi^{(j)}}_{\text{skip-connections}} \right), \quad (6)$$

where $\Phi^{(l)} \in \mathbb{R}^{d_l \times d_{l+1}}$ are the skip-connection parameters.

If the skip-connection is non-parametric and involves only the last layer, i.e., $\Phi^{(l)} \equiv \mathbf{I}_{d_l}$ and $\Phi^{(<l)} \equiv \mathbf{0}$, it is called identity mapping, and it simply “pushes forward” the representation of the nodes at layer l (DHCF [95], HCHA [10]). If parametric, the skip-connection can either have a different set of parameters w.r.t. the convolution (DualHGNN [193], HAIN [15]) or share the same set of parameters, i.e., $\Theta^{(l)} \equiv \Phi^{(l)}$ with $\Phi^{(<l)} \equiv \mathbf{0}$ (SHCN [38], HHNE [19], MGCN [37], HyperRec [175]). Note that if \mathbf{R} has a non-zero diagonal (i.e., it considers self-loops), the skip-connection of the previous layer with shared parameters is already included in Eq. (5) (e.g., HGNN [55]). The identity mapping can also be performed on the aggregated messages rather than the previous node representation(s), and it can be combined with the “standard” skip-connection as in ResHGNN and UniGCNII [90].

A particular case of the skip-connection is the initial residual connection which pushes forward only the initial representation instead of the representation of the previous layer, i.e., $\Phi^{(>0)} \equiv \mathbf{0}$. For instance, in ResHGNN [87], and HyperINF [98], the residual connection is not parametrized but weighted, specifically, $\Phi^0 \equiv \alpha_l \mathbf{I}$ and $\Theta^{(l)} \equiv (1 - \alpha_l) \mathbf{I}$, where α_l is a hyperparameter that balances the contributions of the previous representations. An alternative to the

additive skip-connection is the multiplicative-skip connection that aims to learn a nonlinear gate to modulate the embeddings at a feature-wise granularity through dimension re-weighting (MHCN [206]).

Attention mechanism. The aggregation function in Eq. (5) represents the (normalized) average overall neighbor nodes. Even though straightforward, this choice might be sub-optimal as it gives all neighbors the same importance. A natural way to overcome this limitation is to weigh each neighbor differently. This solution is the basic idea behind the attention mechanism, i.e., a neural network that learns how to weight the neighbor nodes, first introduced for graphs by Velickovic et al. [171]. In a nutshell, (hyper)graph attention is a particular case of (H)GNN in which the message sent to the neighbors is not the same.

The literature offers different types of attention mechanisms, and the most popular in hypergraphs are additive attention (DHCN [190], DHGNN [97], DualHGNN [182], HNN [160], HGC-RNN [202]) and multiplicative attention (HyperGAT [46], SHARE [176], [188], HGNN [47], HGAT [33], DH-HGCN [78], [161]). Although less popular on hypergraphs, it is also possible to add multiple attention “heads” (HCHA [10], DHAT [120], SHT [187]), where each one computes k different attention weights on independently parameterized attention layers. Eventually, the aggregated messages using the k attentions are linearly combined. Attention mechanisms can also help combine node embeddings in tasks like group recommendation (HHGR [211], HCR [96]) or handling temporal sequences (e.g., STHAN-(S)R [153], (D)STHGNN [178]). Recently, thanks to the success of the transformer architecture [170], self-attention mechanisms are starting to appear in convHNN (DHGNN [97], HOT [101], SSF [83], Hyper-SAGNN [217], HyperRec [175]), as widely used to combine sequential node/hyperedge embeddings for task-specific purposes (Higashi [216], DualHGNN [182], HyperTeNet [172], DHAT [120], H²SeqRec [110]).

Gated updates. Before the advent of transformer architectures, recurrent neural networks were the standard tool to deal with sequential inputs. In hypergraphs, the sequential inputs can refer to (i) node-level, i.e., node features change over time, or (ii) structural-level, i.e., the hypergraph topology changes over time. In the node-level case, the neighbors’ information contains not only their current hidden state but also a completely new observation. In this setting, the main mechanisms used are Gated Recurrent Units, i.e., GRU ([188], HGC-RNN [202]), Gated GNN [111] ([179]) and Gated Linear Units, i.e., GLU, stacked on top of convolutional layers (MT-HGCN [177], (D)STHGNN [178]). As far as we know, there are no relevant hypergraph embedding techniques dealing with structural sequential input.

Spectral convolution. Spectral hypergraph convolution relies on the spectral (hyper)graph theory, and it is the basis of the very first convHNN method proposed by Feng et al. [55] (HGNN). As previously discussed, HGNN can also be interpreted as a special case of the MPF (see Eq. (5)); however, it could not be straightforward to directly map a particular spectral hypergraph convolution method in the MPF, though it could be argued that is possible [13].

The main concept behind spectral hypergraph convolution is the connection between the Fourier transform and the Laplace operator. We can define hypergraph Fourier transforms of a hypergraph signal through the eigendecomposition of the hypergraph Laplacian, where the eigenvalues represent hypergraph frequencies (i.e., the spectrum of the hypergraph), and eigenvectors denote frequency components (i.e., hypergraph Fourier basis). Then the convolution in the spectral domain is defined via point-wise products in the transformed Fourier space (HpLapGNN [57], pLapHGNN [123]).

Instead of Fourier basis, it is possible to employ Wavelet basis as in HGWNN [134], and HWNN [159]), where the hypergraph wavelet-transform projects graph signal from the vertex domain onto the spectral domain.

5.3.2 Random walk-based approaches. Random walk (RW) is a common alternative to convolution to encode the concept of closeness. A nice and commonly used property of an RW on a (hyper)graph is its analogy with sentences in a text: a

path is a sequence of nodes (and hyperedges) like a phrase is a sequence of words. This similitude allows leveraging the distributional hypothesis in linguistics, stating that words that occur in the same contexts tend to have similar meanings [80]. Likewise, nodes sharing the same structural context should be close to each other in the latent space. This analogy enables the application of natural language processing techniques to (hyper)graphs (DeepWalk [140]).

In this context, an established approach follows the two-steps framework introduced in DeepWalk: first, walks in the hypernetwork capture the structural context, then node embeddings are learned via a natural language processing model, usually Skip-gram and CBOw [129]. Finally, the so-obtained node representations can be either the input of a readout layer (Hyper2Vec/NHNE [85, 86], Hyper-gram [89], DHHE [41], HEMR [105]) or of another neural network module that fine-tunes the embeddings (DHE [138], Hyper-SAGNN [217]). In [225] (HRSC), the just described method (using CBOw) is applied on the two-section graph, which, however, fails to learn the representation of the hyperedges. Thus, to address this issue, the authors proposed a novel negative sampling-based set constraint objective function.

The nature of hypergraphs allows to inherently generalize existing graph RW approaches to these structures, but their richer semantics requires defining hypergraph-specific strategies. For instance, resembling the exploration-exploitation strategy proposed by node2vec [68], Hyper2Vec [86] exploits a degree biased 2nd order RW to sample paths in a homogeneous hypergraph. A similar strategy is also proposed in Hyper-SAGNN [217]. Conversely, the hyper-path-based RW approach of Hyper-gram [89] captures the idea that hyperedges may encode indecomposable relationships in a heterogeneous hypergraph. On the same line, DHE [138] also proposes a new RW model to acquire co-member information in each hyperedge.

5.3.3 Encoder-based approaches. When it comes to learning low-dimensional embeddings, the encoder-decoder architecture is the go-to approach in machine learning. However, this type of architecture may fail in capturing structural information encoded by hypergraph-like structured data. For this reason, these networks are usually included as a sub-module of bigger networks. For instance, Hyper-SAGNN [217] uses this approach to initialize node representations, while DHNE [168] uses an autoencoder to learn a first latent representation for the nodes, then fine-tuned using a neural network specifically designed to preserve the second-order proximity. More “sophisticated” approaches make use of Variational Autoencoders (HeteHG-VAE [52]) and DeepSets (DHE [138]). In general, when coupled with specific decoders, these methods are mostly used in task-specific hypergraph embeddings (e.g., HeteHG-VAE [52], Event2Vec [42], HNN-HM[112]).

Remarks. Thanks to their versatility, HNNs can be easily adapted to a broad spectrum of tasks, such as (multi-class) classification (e.g., HpLapGCN [57], HyperGCN [195], DualHGNN [182], DHE [138]), regression (e.g., [188], HGC-RNN [202]), link prediction (e.g., HeteHG-VAE [52], DHCF [95], DHNE [168]), and recommendation (e.g., HHGR [211], HyperRec [175], SHARE [176]). The modularity of NN methods also eases the fusion of different input channels. For instance, using the dual of a hypergraph makes it possible to have the same architecture designed for hypergraph node embeddings to learn hyperedge embeddings (e.g., NHNE [85]). In general, multi-channel input can be helpful in domain-specific tasks, for example, when dealing with heterogeneous hypergraphs (e.g., MHCN [206], DHCF [95]). HNNs can also manage dynamic hypergraphs [188], or in general, tasks that require modeling sequential information (e.g., H²SeqRec [110]).

Moreover, HNNs allow performing transfer learning and reusing well-established architectures on various tasks by simply designing the readout layer(s) or a task-specific loss function. The typical learning setting of HNNs is (semi-) supervised transductive learning with only very few exceptions like DHE [138], G-MPNN [194], AllSet [40] and HGAT [33] that can be used in inductive settings.

Method	Input HG	Embedding Output	MPF	Skip-conn	Attention	Gated upd	Spectral	RW	Enc-based	Task	Venue	Year	Code (github.com/)
DHNE [168]	Het, U	\mathcal{V}							✓	Class, LP	AAAI	2018	tadpole/DHNE
DHHE [41]	Het, NU	\mathcal{V}						✓		Class, LP	ICME	2018	-
HHNE [19]	Het, NU	\mathcal{V}	✓	✓						LP	ICDM	2018	ilidanlab/HHNE
HGNN [55]	Hom, NU	\mathcal{V}	✓	✓						Class	AAAI	2019	iMoonLab/HGNN
Hyper2Vec [86]	Hom, NU	\mathcal{V}						✓		Class	DASFAA	2019	jeffhj/NHNE
DHGNN [97]	Hom, NU	\mathcal{V}	✓		✓					Class	IJCAI	2019	iMoonLab/DHGNN
HplapGCN [57]	Hom, NU	\mathcal{V}					✓			Class	Neurocom.	2019	-
Hyper-gram [89]	Het, U	\mathcal{V}						✓		LP	CIKM	2019	HKUST-KnowComp/HPHG
HyperGCN [195]	Hom, NU	\mathcal{V}	✓							Class	NeurIPS	2019	mallabisc/HyperGCN
DHE [138]	Hom, NU, MC	\mathcal{V}, \mathcal{E}						✓	✓	Class	NeurIPS	2019	Josh-Payne/deep-hyperedges
Hyper-SAGNN [217]	Het, NU	\mathcal{V}			✓			✓	✓	LP	ICLR	2020	ma-combio/Hyper-SAGNN
HyperRec [175]	Het, NU, D	\mathcal{V}	✓	✓		✓				Recom	SIGIR	2020	-
NHNE [85]	Hom, NU, MC	\mathcal{V}, \mathcal{E}						✓		Class, LP	TOIS	2020	jeffhj/NHNE
HNHN [49]	Hom, NU	\mathcal{V}, \mathcal{E}	✓							Class	ICMLW	2020	-
MGCN [37]	Hom, NU	\mathcal{V}	✓	✓						LP	KDD	2020	-
DHCF [95]	Hom, NU, MC	\mathcal{V}	✓	✓						Recom	KDD	2020	-
HGC-RNN [202]	Hom, NU	\mathcal{V}	✓		✓	✓				Regr	KDD	2020	-
NHP [196]	Het, NU, Dir	\mathcal{V}	✓							LP	CIKM	2020	-
AdaHGNN [183]	Hom, NU	\mathcal{V}	✓							Class	MM	2020	-
HyperGAT [46]	Het, NU, MC	\mathcal{V}, \mathcal{E}	✓		✓					Class	EMNLP	2020	kaize0409/HyperGAT
SHCN [38]	Het, U	\mathcal{V}	✓	✓						Recom	TOIS	2020	-
G-MPNN [194]	Hom, NU	\mathcal{V}	✓							Class, LP	NeurIPS	2020	naganandy/G-MPNN-R
HGAT [33]	Hom, NU	\mathcal{V}	✓		✓					Class	TrustCom	2020	-
HAIN [15]	Hom, NU	\mathcal{V}	✓	✓						Class	IEEE BigData	2020	-
HCR [96]	Het, NU	\mathcal{V}	✓		✓					Recom	ICDM	2021	-
HCHA [10]	Hom, NU	\mathcal{V}	✓	✓	✓					Class	Pat Rec	2021	-
DHCN [190]	Hom, NU	\mathcal{V}	✓		✓					Recom	AAAI	2021	xiaxin1998/DHCN
HWNN [159]	Het, NU, MC	\mathcal{V}	✓				✓			Class	WSDM	2021	-
STHGNN [178]	Hom, NU, D, MC	\mathcal{V}	✓			✓				Regr	TITS	2021	-
HGCELM [115]	Hom, NU	\mathcal{V}	✓							Class	Appl Sci	2021	-
MHCN [206]	Het, NU, MC	\mathcal{V}	✓	✓						Recom	TheWebConf	2021	Coder-Yu/QRec
HNN [160]	Hom, NU	\mathcal{V}	✓		✓					LP	TheWebConf	2021	-
DualHGNN [193]	Hom, NU, MC	\mathcal{V}	✓	✓						Class, LP	TheWebConf	2021	xuehansheng/DualHGNN
KHNN [114]	Hom, NU, MC	\mathcal{V}	✓							Recom	DASFAA	2021	-
SSF [83]	Hom, U	\mathcal{V}	✓		✓					LP	DASFAA	2021	-
STHAN-SR [153]	Hom, U, D	\mathcal{V}	✓		✓					Regr	AAAI	2021	-
SHARE [176]	Hom, NU	\mathcal{V}, \mathcal{E}	✓		✓					Recom	SDM	2021	-
[156]	Hom, NU	\mathcal{V}, \mathcal{E}	✓							LP	SDM	2021	-
MHGNN [9]	Hom, NU	\mathcal{V}	✓							Class	TIP	2021	-
[188]	Het, NU, D	\mathcal{V}	✓		✓	✓				Regr	JIM	2021	-
HeteHG-VAE [52]	Het, NU	\mathcal{V}, \mathcal{E}			✓				✓	LP	PAMI	2021	haoyfan/HeteHG-VAE
UniGNN [90]	Hom, NU	\mathcal{V}, \mathcal{E}	✓	✓						Class	IJCAI	2021	OneForward/UniGNN
[69]	Het, NU	\mathcal{V}	✓		✓					LP	WWW	2021	-
S ² HGN [124]	Hom, NU, MC	\mathcal{V}	✓							Class	SPIE RS	2021	-
pLapHGNN* [123]	Het, NU	\mathcal{V}					✓			Class	TMM	2021	-
DualHGNN [182]	Hom, NU, MC	\mathcal{V}	✓		✓					Class	J Kno Sys	2021	-
(Res Multi)HGNN [87]	Hom, NU	\mathcal{V}	✓	✓						Class	ICIP	2021	OneForward/ResMHGNN
HNN-HM [112]	Hom, U, Dir	\mathcal{V}	✓		✓				✓	Class	ICCV	2021	-
H ² SeqRec [110]	Het, NU, MC, D	\mathcal{V}	✓							Recom	CIKM	2021	Abigale001/h2seqrec
HHGR [211]	Hom, NU	\mathcal{V}	✓		✓					Recom	CIKM	2021	0411tony/HHGR
HyperCTR [81]	Hom, NU, D	\mathcal{V}	✓							Recom	CIKM	2021	-
HybridHGNN [88]	Hom, NU, MC	\mathcal{V}	✓							Class	PRVC	2021	-
HGDD [136]	Het, NU	\mathcal{V}	✓							Class, LP	ISBRA	2021	-
HGNNa [47]	Hom, NU	\mathcal{V}	✓		✓					Recom	JPCS	2021	-
HGWNN [134]	Hom, NU	\mathcal{V}					✓			Class	Neurocom.	2021	-
EHGNN [99]	Hom, NU	\mathcal{V}, \mathcal{E}	✓							Class	NeurIPS	2021	harryjo97/EHGNN
HOT [101]	Hom, U	\mathcal{V}	✓		✓					LP	NeurIPS	2021	hw9730/hot
HyperTeNet [172]	Het, U, MC	\mathcal{V}	✓		✓					Recom	ICDM	2021	mvijaikumar/HyperTeNet
HyperGroup [73]	Hom, NU	\mathcal{E}	✓	✓						Recom	TOIS	2021	-
F ² HNN [125]	Hom, NU, MC	\mathcal{V}	✓							Class	TGRS	2022	-
DHAT [120]	Hom, NU, Dir	\mathcal{V}	✓		✓					Regr	ITS	2022	-
[179]	Het, NU	\mathcal{V}	✓			✓				Recom	WWW	2022	-
HyperINF [98]	Hom, NU, D, MC	\mathcal{V}	✓	✓						Recom	TSUSC	2022	-
HEMR [105]	Het, NU, Dir	\mathcal{V}						✓		Recom	TNNLS	2022	picuslab/HEMR
GC-HGNN [139]	Hom, NU	\mathcal{V}	✓							Recom	J ECRA	2022	-
HRSC [225]	Het, NU	\mathcal{V}, \mathcal{E}						✓		Class, LP	Appl Sci	2022	-
IHGNN [39]	Het, U	\mathcal{V}, \mathcal{E}	✓							Recom	TheWebConf	2022	CDboyOne/IHGNN
AllSet [40]	Hom, NU	\mathcal{V}, \mathcal{E}	✓		✓					Class	ICLR	2022	jianhao2016/AllSet
HyperSaR [166]	Het, NU	\mathcal{V}, \mathcal{E}	✓							Recom	ECIR	2022	naver/hypersar
MT-HGNN [177]	Hom, NU, D, MC	\mathcal{V}	✓			✓				Regr	TITS	2022	-
[161]	Hom, NU, D	\mathcal{V}, \mathcal{E}	✓		✓					Class	TKDE	2022	-
HCCF [186]	Hom, NU	\mathcal{V}	✓	✓						Recom	SIGIR	2022	akaxlh/HCCF
HGNN* [62]	Hom, NU	\mathcal{V}	✓							Class, Recom	PAMI	2022	iMoonLab/DeepHypergraph
SHT [187]	Hom, NU	\mathcal{V}, \mathcal{E}	✓	✓	✓					Recom	KDD	2022	akaxlh/SHT
DH-HGNN [78]	Hom, NU	\mathcal{V}	✓		✓					Recom	SIGIR	2022	-
LE[198]	Hom, NU	\mathcal{V}	✓	✓						Class	CIKM	2022	ycq091044/LEGCN

Table 4. (Deep) Neural network methods. Input HG: Hom/Het stands for homogeneous/heterogeneous, (N)U for (non-)uniform, MC for multi-channel, D for dynamic, and Dir for directed. Task: Class stands for classification, LP for link prediction, Regr for regression, and Recom for Recommendation. All methods can handle attributed nodes.

As for (deep) NNs in general, HNNs require much data to be adequately trained, and the hyper-parameters tuning can be costly. Another drawback of HNNs, especially on very deep architectures, is that the training phase may require high computational power (e.g., GPUs, TPUs), hampering the applicability on edge devices.

Among all the HNNs techniques, convHNNs are nowadays predominant in the literature thanks to their superior performance. This trend is evident from Table 4: both encoder-based and random-walk-based methods have almost disappeared since 2020. The success of such a methodology is mainly due to the reuse of the knowledge coming from the graph representation learning literature. Akin to GNNs, these approaches learn the node representations that can then be aggregated to learn the latent representations of hyperedges (e.g., NHP [196]) or the whole hypergraph (e.g., [188]). Recently the trend seems to be drifting towards convolutional methods specifically designed for HGs (e.g., [156], AllSet [40]), where the hyperedge embeddings are explicitly learned.

However, convHNNs are generally more demanding in terms of computational requirements than encoder-based and random-walk-based methods; thus, such approaches should be preferred in case of limited computational capabilities. Further, the clique-expansion transformation used by many convolutional approaches is a possible limitation since the high-order nature of the relations is only indirectly exploited in the learning process. An important strength of random-walk-based methods is the universality of the approach that can be (almost) seamlessly applied to any hypergraphs. Encoder-based methods, instead, may require some effort in the design of the encoder/decoder architecture. Still, when combined together, random-walk and encoder-based methods can achieve state-of-the-art performance, e.g., Hyper-SAGNN [217] and DHE [138].

Task-wise, we do not notice any particular trend: all (deep) neural network approaches have been applied successfully on classification, link prediction, and recommendation tasks. Nonetheless, when dealing with dynamic or multi-channel hypergraphs, convHNNs are preferable since they can easily be designed to handle recurrent or multi-channel inputs.

5.4 Comparison of Hypergraph Representation Learning Models

All methods discussed in the previous sections have distinctive characteristics (see Table 5), making them fit specific application domains and constraints. In general, we can observe that *spectral-based* embedding algorithms have been little used in real-world, large-scale applications due to their scalability issues. Nonetheless, these methods represent a viable choice when dealing with small datasets. The simple design of *proximity-preserving* approaches makes them particularly suitable when we aim to take advantage of prior domain knowledge that can easily embed in the learning procedure via the design of specific similarity functions. Akin to spectral methods, PP methods work better on small to medium-scale data sets. Currently, (*Deep*)NN methods embody the state-of-the-art in hypergraph representation learning, and the vast majority of embedding methods belong to this category. This trend directly derives from the high-quality performance these methods achieve and the possibility of performing end-to-end learning. Further, the flexibility of the MPF allows capturing the high-order relationships encoded by each hyperedge. Hypergraph NNs scale better than the other families of HG embedding techniques and, thus, can be applied to large-scale data sets. However, the training of particularly deep networks may be hardware-demanding.

6 LEARNING TASKS AND APPLICATIONS

Hypergraph representation learning is functional to solve a plethora of graph analytics tasks as the learned latent representations allow the efficient execution (in both time and space) of traditional vector-based algorithms. Based on consolidated surveys on graph representation learning [16, 29], in this work, we classify existing applications in

Embedding Method	Pros	Cons
<i>Spectral learning</i>	(i) Consider global structure; (ii) Solid theoretical foundation; (iii) Easy to implement and parallelize.	(i) Cannot capture high-order interactions (transformation to clique graphs); (ii) Scalability issues (in terms of computational time and storage).
<i>Proximity-preserving</i>	(i) Simple and versatile design based on node similarity maximization; (ii) Easy to adapt to specific tasks.	(i) May not model high-order non-linear relations beyond single-hop neighbors (transformation to clique graphs); (ii) May suffer from scalability issues.
<i>(Deep) Neural networks</i>	(i) Capture high-order and non-linear relations; (ii) High versatility; (iii) Can easily embed both structural and additional information; (iv) Transfer learning capabilities.	(i) Hardware demanding training; (ii) High-computational cost; (iii) May require an extensive phase of hyper-parameters tuning.

Table 5. Comparison of hypergraph representation learning techniques.

which hypergraph embedding came into play according to whether they focused on nodes or hyperedges. It is worth noting that we only match each work with the task explicitly reported in the corresponding article. Table 6 lists the most commonly used datasets across the works reviewed. For each dataset, we detail whether it belongs to a specific category, the link where it can be downloaded, a reference to the paper introducing it (if any), which methods have been tested on it and for which tasks, and how it has been modeled with a hypergraph. For the sake of brevity, we only report information about the datasets used by at least two articles and whose data are currently accessible. A complete list can be found at the following link <https://tinyurl.com/hgrl-dataset>.

6.1 Node-related learning tasks and applications

In this section, we focus on describing the role of hypergraph representation learning in relation to classical learning tasks defined on the nodes of a hypernetwork, e.g., classification, clustering, and recommendation. For each task, we further describe the application domain(s) where it has been exploited.

6.1.1 Node classification. Node classification is a supervised task in which the goal is to assign the correct label to each unlabeled node in the hypergraph based on the patterns learned from the other already labeled nodes (supervision). Node classification is one of the most common tasks discussed in the hypergraph embedding literature and finds wide application in (possibly multi-label) event (HEBE-(PE/PO) [71, 72], Event2Vec [42]), movie (e.g., DHNE [168], HRSC [225], DualHGNN [182]), authorship (e.g., HyperGCN [195], Hyper2Vec/NHNE [85, 86], HWNN [159]), citation (e.g., DHGNN [97], HCHA [10], UNIGNN [90], LE [198]), image (e.g., SSHGDA [118], AdaHGNN [183], HGWNN [134]), and item (e.g., DualHGNN [193], AllSet [40], DualHGNN [193]) classification, among others (e.g., HOT [101], EHGNN [99], DHE [138]).

Node classification is applied in a (semi-)supervised setting, and, in the literature, there are two main approaches to tackle it. The first approach consists of a three-step sequential procedure, in which (i) the embedding method first learns the latent representation of the nodes, (ii) then an existing classifier is trained on the labeled instances, and (iii) finally, the trained classifier predicts the class of an unlabeled node given its embedding. SVM (e.g., DHNE [168], HEBE-(PE/PO) [71, 72], SSHGDA [118]), logistic regression (e.g., HEBE-(PE/PO) [71, 72], Hyper2Vec/NHNE [85, 86], Event2Vec [42]), and k-nearest neighbors (e.g., SSHGDA [118]) are the most commonly used classifiers. The second approach relates to (deep) neural networks. Specifically, the task is solved in an end-to-end fashion where the first part of the network explicitly or implicitly learns the node embeddings, while the readout layer(s) solve the specific classification (e.g., AllSet [40], HOT [101], HNHN [49]).

Dataset type	Dataset	Link (https://tinyurl.com/)	Classification	Clustering	Link Prediction	Recommendation	Input setting	Used by
<i>Categorical</i>	20newsgroup [100]	20newsgroup	✓	✓			Hom, NU	[10, 40, 141, 150, 175, 198]
	Cancer [157]	breast-cancer-db	✓	✓			Hom, NU	[115, 150]
	Zoo [56]	zoo-db	✓	✓			Hom, NU	[40, 141, 150, 198, 221]
<i>Co-authorship/ Co-citation</i>	DBLP* [61, 201]	dblp-net dblp-bip aminer-dblp	✓		✓	✓	Hom/Het, NU, MC	[15, 40, 49, 52, 69, 71, 72, 85, 86, 90, 159, 160, 164, 182, 195, 196]
	Cora* [154]	cora-cit-network-db cora-coauth	✓		✓		Hom/Het, NU, MC	[10, 15, 40, 55, 62, 69, 88, 90, 97, 138, 141, 159, 160, 182, 195, 198]
	Citeseer [147]	citeseer	✓				Hom, NU, MC	[10, 15, 40, 49, 57, 62, 88, 90, 97, 141, 182, 195, 198]
	PubMed* [154]	pubmed-db pb-diabetes	✓				Hom/Het, NU, MC	[10, 15, 40, 49, 55, 62, 88, 90, 138, 159, 195, 198]
<i>Images</i>	Botswana	botswana-img	✓				Hom, NU, A, MC	[119, 125, 162, 207]
	Houses	house-img	✓	✓			Hom, (N)U, Dir	[112, 142]
	Indian Pines [18]	indian-pines-db	✓				Hom, NU, A, MC	[118, 124, 125, 162, 207]
	KSC	ksc-db-img	✓				Hom, NU, A, MC	[119, 124, 125]
	ModelNet40 [185]	modelnet40	✓				Hom/Het, NU, MC	[9, 40, 55, 62, 87, 88, 123, 134, 198]
	NUS-WIDE [43]	nus-wide-db	✓				Het, NU	[41, 183]
	NTU [†] [35]	http://3d.csie.ntu.edu.tw	✓			✓	Hom/Het, NU, MC	[9, 40, 55, 62, 87, 88, 123, 134, 198]
	Pascal VOC [50]	pascal-voc-2007	✓			✓	Hom, (N)U, Dir	[112, 183]
Pavia University	pavia-uni	✓				Hom, NU, A, MC	[118, 125, 162]	
<i>Miscellanea</i>	CiteULike [70, 173]	citeulike-db				✓	Hom/Het, NU, MC	[95, 224]
	Drug	drug-faers	✓			✓	Het, (N)U	[19, 89, 101, 168, 217, 225]
	GPS [220]	gps-ms-data	✓			✓	Hom/Het, (N)U	[89, 101, 168, 217, 225]
	wordnet [27, 130]	wordnet	✓			✓	Het, (N)U	[89, 168, 217, 225]
<i>Movies</i>	CAMRa2011*	CAMRa2011	✓			✓	Hom/Het, NU	[96, 211]
	IMDB*	imdb-movie-db	✓			✓	Hom/Het, NU, MC	[52, 85, 86, 99, 182]
	MovieLens* [79]	movie-lens	✓			✓	Hom/Het, (N)U, MC, D	[62, 81, 89, 95, 101, 114, 166, 168, 186, 217, 225]
<i>Purchase</i>	Amazon* [107, 133, 201]	amazon-meta amzn-prod amzn-rev	✓		✓	✓	Hom/Het, (N)U, D, MC	[38, 39, 110, 141, 164, 175, 186, 193]
	Diginetica	diginetica				✓	Hom/Het, (N)U	[39, 47, 139, 176, 190]
	Tmall	tmall-ijcai15				✓	Hom/Het, NU	[139, 179, 187, 196]
	YooChoose [22]	yoochoose				✓	Hom, U, D	[47, 176]
<i>Review</i>	Goodreads* [174, 175]	goodreads-db				✓	Het, NU, D	[110, 172, 175]
	Yelp*	yelp-db	✓			✓	Hom/Het, NU	[38, 40, 52, 69, 71-73, 78, 160, 186, 187, 206]
<i>Social networks</i>	Foursquare [199]	foursquare-net				✓	Het, U	[167, 199, 200]
	Douban	douban-db				✓	Hom/Het, (N)U, MC	[52, 73, 78, 83, 206, 211]
<i>Songs</i>	#nowplaying [209]	nowplaying-db				✓	Hom, NU	[139, 190]
	Last.FM	last-fm-db				✓	Hom/Het, NU, MC	[114, 164, 166, 206]
	Spotify [24]	million-songs spotify-db				✓	Het, NU, Dir	[105, 172]

Table 6. Most commonly used public datasets by the articles reviewed. Input setting: Hom/Het stands for homogeneous/heterogeneous, (N)U for (non-)uniform, MC for multi-channel, D for dynamic, and Dir for directed. The symbol * means that the specific dataset is available in multiple versions. [†] The NTU dataset is not reachable from our location. All links were accessed on the 12th of June, 2023.

6.1.2 Node clustering. The goal of node clustering is to group similar nodes together in a way that nodes in the same cluster are more similar to each other than nodes in other groups. Node clustering is an example of an unsupervised learning setting, and it is usually applied when no node labels are available to find affinities between groups of nodes. The classical workflow is first to learn the low-dimensional representations of the nodes to then apply the desired traditional clustering algorithm on the latent vectors. All works included in this survey (Zhou et al. [221], Ren et al. [142], Saito et al. [150], Event2Vec [42]) adopt k-means as clustering algorithm. Specifically, Event2Vec [42] exploits this task to test the expressiveness of the learned node embeddings in capturing within-event and cross-event relationships in event data, while Ren et al. address the problem of clustering object views.

6.1.3 Node Recommendation. The task of node recommendation consists in finding the best nodes of interest (e.g., items) for a given node (e.g., user) based on certain criteria [51, 145]. In this context, hypergraphs embody a valuable tool to structurally encode the inherent high-order relationships that may arise beyond the classical user-item relation. For instance, hypergraphs can explicitly model additional information attached to the user-item interaction, such as emotions (e.g., HGE [204], SHCN [38]), item categories or properties (e.g., HHE [224], HEMR [105], HyperCTR [81]), search query (e.g., HyperSAR [166], IHGNN [39]), purchase history (e.g., HyperREC [175], H²SeqRec [110]), social relationships (e.g., MHCN [206]), user sessions (e.g., SHARE [176], DHCN [190], GC-HGNN [139]), or user/item groups (e.g., DHCF [95], HHGR [211], HCR [96]). In general, recommender systems based on hypergraphs have been used for document/book (e.g., HHE [224], H²SeqRec [110], KHNN [114]), movie (e.g., HGE [204], DHCF [95], HHGR [211]), product (e.g., HyperRec [175], SHCN [38], SHARE [176]), song (e.g., FOBE/HOBE [164], MHCN [206], HEMR [105]), video (e.g., HyperCTR [81]), business/place (e.g., HyperSAR [166], MHCN [206], HHGR [211], DH-HGCN [78]), news (e.g., DHCF [95]) recommendation, and others (e.g., GC-HGNN [139], HCR [96], [179]).

6.2 Hyperedge related tasks and applications

Next, we discuss the role of hypergraph representation learning concerning the link prediction task, its variations, and associated application domain(s).

6.2.1 Link prediction. In traditional graphs, the link prediction task involves inferring new relationships or still unknown interactions between pairs of entities based on their properties and the currently observed links [113]. In hypergraphs, the problem of hyperlink prediction involves predicting missing hyperedges from the set $2^{|V|} \setminus \mathcal{E}$ based on the current set of observed hyperedges \mathcal{E} . When dealing with k -uniform hypergraphs, the size of the hyperedge to predict is bounded by the nature of the modeled relation [101, 168, 217]. However, in the more general setting of non-uniform hypergraphs, the variable cardinality of a hyperedge makes link prediction methods defined for graphs infeasible as they are based on exactly two input features, i.e., those of the two vertices potentially forming a link [128, 196].

To overcome the difficulties mentioned above, all methods devised for hyperlink prediction heavily rely on negative sampling techniques to discard not meaningful relations in both training and evaluation phases (e.g., MSC-LBSN [167], NHP [196]). Then, the existence of a given hyperedge is based on a similarity score evaluated over pairs of embedding vectors corresponding to the nodes that should be comprised in the relation. The similarity score can be either computed with some similarity measures, such as euclidean distance and cosine similarity (e.g., NHNE [85], LBSN2Vec [199, 200]), or a more complex edge classifier, such as logistic regression. In this last case, the link prediction task is treated as a classification problem where the target class label indicates the presence or absence of a link between a pair of nodes (e.g., G-MPNN [194], HNN [160], DualHGNC [193]).

In the literature, this task has been applied to detect ⟨user, location, activity⟩, ⟨user, movie, tag⟩, ⟨user, drug, reaction⟩, and ⟨synset, relation type, synset⟩ relations (e.g., DHNE [168], Hyper-SAGNN [217], HRSC [225], HOT [101]). Other application domains relate to predicting hyperlinks in non-uniform (possibly) heterogeneous collaboration (e.g., HeteHG-VAE [52], NHP [196], NHNE [85]), user-movie interest (e.g., HeteHG-VAE [52], NHNE [85]), user-item (e.g., DualHGNC [193], [179], HNN [160]), and chemical/drug reaction (e.g., HHNE [19], HGDD [136]) networks.

6.2.2 Network reconstruction. The task of network reconstruction can be seen as a particular case of link prediction, in which all hyperedges of the original hypernetwork need to be inferred (DHNE [168], Hyper-gram [89]). Another variation of the problem, a.k.a. hyperedge expansion, has been proposed by Srinivasan et al. [156] and is based on set theory. In this case, the task consists of predicting the missing nodes (elements) from a given hyperedge (set).

6.3 Other applications

In the following, we provide a brief description of other relevant applications.

6.3.1 Time-series forecasting. Time series forecasting is a subclass of regression problems, and, specifically, it is the task of fitting a model to historical, time-stamped data in order to predict future values. This umbrella term includes several domain-specific tasks in which hypergraphs have been exploited, such as traffic prediction (MT-HGCN [177], DHAT [120]), passenger flow prediction (STHGCN [178]), gas and taxi demand (HGC-RNN [202]), and stock selection (STHAN-SR [153]). In this context, homogeneous and non-uniform hypergraphs model spatio-temporal information.

6.3.2 Visualization. Visualizing the outcomes of a hypergraph representation learning algorithm embodies a strong demonstration of whether the devised embedding method is preserving the desired characteristic of the input network. After learning the low-dimensional latent vectors, those are further projected into a two-dimensional space, usually via t-SNE [169], the state-of-the-art tool to visualize high-dimensional data. Another way to go, common for spectral-based methods, is to plot each node representation considering the two or three smallest eigenvectors of the Laplacian matrix [221]. In the case of node classification or, when possible, clustering tasks, each node category is colored differently: in this way, it is easy to grasp whether nodes belonging to the same category or sharing the same characteristics are embedded close to each other. Some methods discussed in this survey (e.g., Event2Vec [42], HyperGCN [195], DHE [138], HHNE [19]) use this technique to compare their computed embeddings against some baselines.

6.3.3 Knowledge Hypergraphs. When it comes to model knowledge bases (KBs), graphs are the go-to approach since they naturally represent ternary relations, i.e., facts of the form $\langle head, relation, tail \rangle$, in KBs. Nonetheless, with n-ary relations, hypergraphs could be a valid alternative [53, 197]. Knowledge hypergraph embedding techniques are strictly interlinked to the task (usually predicting facts); hence, it is hard to transfer these methodologies to different contexts.

6.3.4 Natural sciences. Over the last decades, network science has become an established framework to analyze interactions between biological and chemical agents [17]. In this context, the task of hypergraph representation learning has been leveraged to predict the material removal rate in chemical-mechanical processes [188, 189], multi-way chromatic interactions [214, 215], genome features [216], and drug-disease association [136]. Other work also focused on medical-related issues, such as cancer tissue classification [12], autism diagnosis [126], and plant diseases detection [152].

7 FUTURE DIRECTIONS

In the following, we discuss challenges and potential future directions related to hypergraph representation learning.

Deep hypergraph representation learning. The last years have signed the rise of (deep) neural hypergraph embeddings. This trend follows the trail of GNNs, which have shown outstanding performance on graphs. However, most of the efforts have been devoted to adapting existing ideas on graphs to hypergraphs. A real challenge for the future is developing neural models exclusively designed for hypergraphs that can exploit their peculiar characteristics. For instance, most convHNN methods approximate each hyperedge with a clique and, hence, lose the relation’s indecomposable nature. A possible direction is to work toward MP functions that aggregate neighbors’ representations by directly exploiting high-order interactions rather than pairwise connections, following the footsteps of [40, 156].

Beyond node embeddings. Most of the methods presented in this survey deal with node embeddings, with only a few exceptions which also consider hyperedges. Nonetheless, there is a need to develop solutions able to learn

representations for whole hypergraphs or sub-structures. These methods can be useful in visualization [11], community detection [31], and hyper knowledge-graph embedding [148] tasks. Following the spirit of the message-passing framework, a promising direction is learning permutation invariant functions able to aggregate the node/hyperedge representations in a meaningful way for the task at hand.

Dynamic hypergraph embedding. Dynamicity is an essential characteristic of many networks and can manifest either as time-varying features or structural interaction patterns [5, 16]. (Hyper)graphs may evolve in terms of node/edge structure (addition/removal) or information. These characteristics cause the static embedding approach to fail as the hypergraph size is not fixed, features may drift over time, and adding new nodes/edges requires efficiently updating representations [36]. Current work on dynamic hypergraph embedding assumes a transductive setting, with a fixed node set and a variable set of hyperedges [98] or node features [110, 175, 177, 188, 202]. Still, a more challenging problem is predicting newly added nodes' representations. As inductive frameworks exist, the major challenge here is how to update the existing node representations incrementally and adapt to the concept drift of the graph structure. Inspired by graph representation learning literature [76], a promising direction could be to learn functions that generate embeddings by aggregating features from nodes' neighborhoods instead of training individual node embeddings.

Interpretability. Most state-of-the-art hypergraph embedding methods are built using convolutions, usually in a layered architecture. The complexity of such models makes them expressive enough to extract a good condensed hypergraph representation. Yet, the highly non-linear nature of these models harms their interpretability. In general, to interpret an embedding, we need to find an association between latent features and features of the original hypergraph. A possible direction could be exploring the so-called disentangled representations [23]. Disentangled representation learning [74, 109] may help to learn uncorrelated latent features, which may allow characterizing the various underlying explanatory factors behind the learned factorized representations.

Scalability. Scalability is a critical requirement when dealing with many real-world networks. Hypergraphs add further complexity as not only the number of nodes and hyperedges can scale to millions, but also the same hyperedges cardinality may be huge. Currently, little work exists to improve hypergraph algorithms' scalability (e.g., [195]); thus, scalable computational paradigms and models are a critical challenge for future work. NetVec [127] embodies an effort toward this direction by applying a coarsening strategy to preprocess data. In general, hypergraph coarsening or partitioning strategies combined with parallel approaches may represent a valid approach.

Reproducibility. How to directly measure the quality of data representation is a long-standing problem in representation learning [108]. Currently, the quality of the learned representation is measured by indirectly executing a given set of relevant tasks. However, no standard benchmark is currently available in terms of data and methods. In addition to the (not so rare) unavailability of open-source implementations, this fragmentation leads to approaches difficult to reuse and compare. As suggested by Hamilton et al. [77] in the case of graph representation learning, effort should be put into defining a common framework describing - for a given task - the expected network structure to encode, how we expect the models to encode this information, and possible constraints on the learned representation to clarify which approach should be used when in real-world applications.

8 CONCLUSION

The ever-increasing development of computational capabilities and the advancements of theoretical models demand structures to unify and abstract real-world data with complex relations. Hypergraphs are a promising answer to this

need; yet, the overall hypergraph research and, specifically, the hypergraph representation learning field is still in its infancy, but the body of work on the topic is growing rather quickly.

In this survey, we presented a systematic and thorough discussion about hypergraph representation learning. We identified in the literature three main families of approaches to the hypergraph embedding problem: spectral, proximity-preserving, and (deep) neural network hypergraph embedding methods.

- Spectral hypergraph embedding methods represent the first efforts toward learning vectorial representations for hypergraphs. Despite their solid theoretical foundation, spectral methods struggle with large-scale hypergraphs, thus making them less appealing to the research community in recent years.
- Both proximity-preserving and (deep) neural network methods overcome this scalability issue offering two very different approaches to the hypergraph representation learning problem.
 - Proximity-preserving methods use a standard machine learning pipeline where the proximity of the nodes in the latent space is preserved through a loss function defined on the distance (or similarity) between nodes. The loss function can also be designed to embed additional constraints or information tied to the specific task at hand. Such an easy design makes proximity-preserving methods adaptable to more disparate contexts.
 - Nowadays, deep learning-based methods are predominant thanks to their flexibility, capacity, and performance. They take advantage of the experience from the graph representation learning literature and the ever-increasing availability of high computational power that allows training very large (i.e., deep) models. On the one hand, this inheritance from graph representation learning has sped up the development of learning techniques for hypergraphs. On the other hand, it has biased the researchers' efforts toward methods for hypergraphs that are adaptations of the ones for graphs. However, although similar, hypergraphs have inherent characteristics that must be considered, e.g., they encode higher-order relationships. Works like Chien et al. [40] and Srinivasan et al. [156] have opened up new avenues thanks to message-passing functions working on (multi-)sets rather than graph structures. We believe such contributions may inspire the research community to develop novel representation learning frameworks designed explicitly for hypergraphs.

One of the biggest challenges in hypergraph representation learning is the lack of (at least) an affirmed tool that researchers could use as a reference to implement new embedding techniques; while, for the graph counterpart, the community offers many alternatives like DGL⁴, Spektral⁵, and Karate Club⁶, just to name a few. Unluckily, there is no equivalent for hypergraphs, and even libraries for hypernetwork analysis are not so numerous (e.g., HyperNetX⁷[2], SimpleHypergraph.jl⁸[3, 4], Hypergraphx⁹[116]) and as mature as, for example, NetworkX¹⁰ for graphs.

With this survey, we aim to provide future researchers and practitioners with a comprehensive understanding of current trends in addressing the hypergraph embedding problem and sketching some guidelines for designing novel and practical hypergraph representation learning techniques.

⁴<https://docs.dgl.ai>

⁵<https://graphneural.network>

⁶<https://karateclub.readthedocs.io>

⁷<https://pnnl.github.io/HyperNetX>

⁸<https://github.com/pszufe/SimpleHypergraphs.jl>

⁹<https://github.com/HGX-Team/hypergraphx>

¹⁰<https://networkx.org/>

ACKNOWLEDGMENTS

This work has been partially supported by the spoke “FutureHPC & BigData” of the ICSC – Centro Nazionale di Ricerca in High-Performance Computing, Big Data and Quantum Computing funded by European Union – NextGenerationEU, University of Macau (SRG2021-00002-IOTSC) and The Science and Technology Development Fund, Macau SAR (0047/2022/A1).

REFERENCES

- [1] S. Agarwal, K. Branson, and S. Belongie. 2006. Higher Order Learning with Graphs. In *Proceedings of the 23rd International Conference on Machine Learning (ICML '06)*. Association for Computing Machinery, New York, NY, USA, 17–24. <https://doi.org/10.1145/1143844.1143847>
- [2] S. G. Aksoy, C. Joslyn, C. Ortiz Marrero, B. Praggastis, and E. Purvine. 2020. Hypernetwork science via high-order hypergraph walks. *EPJ Data Science* 9, 1 (2020), 16. <https://doi.org/10.1140/epjds/s13688-020-00231-0>
- [3] A. Antelmi, G. Cordasco, B. Kamiński, P. Pralat, V. Scarano, C. Spagnuolo, and P. Szufel. 2019. SimpleHypergraphs.jl—Novel Software Framework for Modelling and Analysis of Hypergraphs. In *Algorithms and Models for the Web Graph*. Springer International Publishing, Cham, 115–129. https://doi.org/10.1007/978-3-030-25070-6_9
- [4] A. Antelmi, G. Cordasco, B. Kamiński, P. Pralat, V. Scarano, C. Spagnuolo, and P. Szufel. 2020. Analyzing, Exploring, and Visualizing Complex Networks via Hypergraphs using SimpleHypergraphs.jl. *Internet Mathematics* 1 (2020), 1–32. Issue 1. <https://doi.org/10.24166/im.01.2020>
- [5] A. Antelmi, G. Cordasco, C. Spagnuolo, and V. Scarano. 2020. A Design-Methodology for Epidemic Dynamics via Time-Varying Hypergraphs. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '20)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 61–69. <https://doi.org/10.5555/3398761.3398774>
- [6] A. Antelmi, G. Cordasco, C. Spagnuolo, and P. Szufel. 2021. Social Influence Maximization in Hypergraphs. *Entropy* 23, 7 (2021). <https://doi.org/10.3390/e23070796>
- [7] D. Arya, D. K. Gupta, S. Rudinac, and M. Worring. 2020. HyperSAGE: Generalizing Inductive Representation Learning on Hypergraphs. <https://doi.org/10.48550/ARXIV.2010.04558>
- [8] G. Ausiello, P. G. Franciosa, and D. Frigioni. 2001. Directed Hypergraphs: Problems, Algorithmic Results, and a Novel Incremental Approach. In *Proceedings of the 7th Italian Conference on Theoretical Computer Science (ICTCS '01)*. Springer-Verlag, Berlin, Heidelberg, 312–327. <https://doi.org/10.5555/646293.687242>
- [9] J. Bai, B. Gong, Y. Zhao, F. Lei, C. Yan, and Y. Gao. 2021. Multi-Scale Representation Learning on Hypergraph for 3D Shape Retrieval and Recognition. *IEEE Transactions on Image Processing* 30 (2021), 5327–5338. <https://doi.org/10.1109/TIP.2021.3082765>
- [10] S. Bai, F. Zhang, and P.H.S. Torr. 2021. Hypergraph convolution and hypergraph attention. *Pattern Recognition* 110 (2021), 107637. <https://doi.org/10.1016/j.patcog.2020.107637>
- [11] Y. Bai, H. Ding, Y. Qiao, A. Marinovic, K. Gu, T. Chen, Y. Sun, and W. Wang. 2019. Unsupervised Inductive Graph-Level Representation Learning via Graph-Graph Proximity. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-19)*. International Joint Conferences on Artificial Intelligence Organization, 1988–1994. <https://doi.org/10.24963/ijcai.2019/275>
- [12] A. B. Bakht, S. Javed, H. AlMarzouqi, A. Khandoker, and N. Werghi. 2021. Colorectal Cancer Tissue Classification Using Semi-Supervised Hypergraph Convolutional Network. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*. 1306–1309. <https://doi.org/10.1109/ISBI48211.2021.9434036>
- [13] M. Balcilar, G. Renton, P. Héroux, B. Gaüzère, S. Adam, and P. Honeine. 2020. Bridging the Gap Between Spectral and Spatial Domains in Graph Neural Networks. <https://doi.org/10.48550/ARXIV.2003.11702>
- [14] M. Balcilar, G. Renton, P. Héroux, B. Gaüzère, S. Adam, and P. Honeine. 2020. When Spectral Domain Meets Spatial Domain in Graph Neural Networks. In *Thirty-seventh International Conference on Machine Learning (ICML 2020) - Workshop on Graph Representation Learning and Beyond (GRL+ 2020)*. 1–9.
- [15] S. Bandyopadhyay, K. Das, and M. N. Murty. 2020. Hypergraph Attention Isomorphism Network by Learning Line Graph Expansion. In *2020 IEEE International Conference on Big Data (Big Data)*. 669–678. <https://doi.org/10.1109/BigData50022.2020.9378335>
- [16] C. D. T. Barros, M. R. F. Mendonça, A. B. Vieira, and A. Ziviani. 2021. A Survey on Embedding Dynamic Graphs. *Comput. Surveys* 55, 1, Article 10 (2021), 37 pages. <https://doi.org/10.1145/3483595>
- [17] F. Battiston, G. Cencetti, I. Iacopini, V. Latora, M. Lucas, A. Patania, J.-G. Young, and G. Petri. 2020. Networks beyond pairwise interactions: Structure and dynamics. *Physics Reports* 874 (2020), 1–92. <https://doi.org/10.1016/j.physrep.2020.05.004>
- [18] M.F. Baumgardner, L.L. Biehl, and D.A. Landgrebe. 2015. 220 Band AVIRIS Hyperspectral Image Data Set: June 12, 1992 Indian Pine Test Site 3. <https://doi.org/doi:10.4231/R7RX991C>
- [19] I. M. Baytas, C. Xiao, F. Wang, A. K. Jain, and J. Zhou. 2018. Heterogeneous Hyper-Network Embedding. In *2018 IEEE International Conference on Data Mining (ICDM)*. 875–880. <https://doi.org/10.1109/ICDM.2018.00104>
- [20] N.C. Behague, A. Bonato, M.A. Huggan, R. Malik, and T.G. Marbach. 2021. The iterated local transitivity model for hypergraphs. *CoRR* abs/2101.12560 (2021).

- [21] M. Belkin and P. Niyogi. 2003. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. *Neural Computation* 15, 6 (2003), 1373–1396. <https://doi.org/10.1162/089976603321780317>
- [22] D. Ben-Shimon, A. Tsikinovsky, M. Friedmann, B. Shapira, L. Rokach, and J. Hoerle. 2015. RecSys Challenge 2015 and the YOOCHOOSE Dataset. In *Proceedings of the 9th ACM Conference on Recommender Systems (RecSys '15)*. Association for Computing Machinery, New York, NY, USA, 357–358. <https://doi.org/10.1145/2792838.2798723>
- [23] Y. Bengio, A. Courville, and P. Vincent. 2013. Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 8 (2013), 1798–1828. <https://doi.org/10.1109/TPAMI.2013.50>
- [24] T. Bertin-Mahieux, D. P.W. Ellis, B. Whitman, and P. Lamere. 2011. The Million Song Dataset. In *Proceedings of the 12th International Conference on Music Information Retrieval (ISMIR 2011)*.
- [25] Á. Bodó, G. Y. Katona, and Pé. L. Simon. 2016. SIS Epidemic Propagation on Hypergraphs. *Bulletin of Mathematical Biology* 78 (2016), 713–735. Issue 4. <https://doi.org/10.1007/s11538-016-0158-0>
- [26] M. Bolla. 1993. Spectra, Euclidean representations and clusterings of hypergraphs. *Discrete Mathematics* 117 (1993), 19–39. Issue 1-3. [https://doi.org/10.1016/0012-365X\(93\)90322-K](https://doi.org/10.1016/0012-365X(93)90322-K)
- [27] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko. 2013. Translating Embeddings for Modeling Multi-relational Data. In *Advances in Neural Information Processing Systems*, C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger (Eds.), Vol. 26. Curran Associates, Inc.
- [28] A. Bretto. 2013. *Hypergraph Theory: An Introduction*. Springer International Publishing, New York, NY (USA). <https://doi.org/10.1007/978-3-319-00080-0>
- [29] D. Cai, X. He, and J. Han. 2007. Spectral Regression: A Unified Subspace Learning Framework for Content-Based Image Retrieval. In *Proceedings of the 15th ACM International Conference on Multimedia (MM '07)*. Association for Computing Machinery, New York, NY, USA, 403–412. <https://doi.org/10.1145/1291233.1291329>
- [30] H. Cai, V. W. Zheng, and K. C. Chang. 2018. A Comprehensive Survey of Graph Embedding: Problems, Techniques, and Applications. *IEEE Transactions on Knowledge and Data Engineering* 30, 9 (2018), 1616–1637. <https://doi.org/10.1109/TKDE.2018.2807452>
- [31] S. Cavallari, V. W. Zheng, H. Cai, K. C.-C. Chang, and E. Cambria. 2017. Learning Community Embedding with Community Detection and Node Embedding on Graphs. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (CIKM '17)*. Association for Computing Machinery, New York, NY, USA, 377–386. <https://doi.org/10.1145/3132847.3132925>
- [32] T. H. H. Chan, A. Louis, Z. G. Tang, and C. Zhang. 2018. Spectral Properties of Hypergraph Laplacian and Approximation Algorithms. *J. ACM* 65, 3, Article 15 (2018), 48 pages. <https://doi.org/10.1145/3178123>
- [33] C. Chen, Z. Cheng, Z. Li, and M. Wang. 2020. Hypergraph Attention Networks. In *2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*. 1560–1565. <https://doi.org/10.1109/TrustCom50675.2020.00215>
- [34] D. Chen, Y. Lin, W. Li, P. Li, J. Zhou, and X. Sun. 2020. Measuring and Relieving the Over-Smoothing Problem for Graph Neural Networks from the Topological View. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 04 (2020), 3438–3445. <https://doi.org/10.1609/aaai.v34i04.5747>
- [35] D.-Y. Chen, X.-P. Tian, Y.-T. Shen, and M. Ouhyoung. 2003. On Visual Similarity Based 3D Model Retrieval. *Computer Graphics Forum* 22, 3 (2003), 223–232. <https://doi.org/10.1111/1467-8659.00669>
- [36] F. Chen, Y. Wang, B. Wang, and C.-C. J. Kuo. 2020. Graph representation learning: a survey. *APSIPA Transactions on Signal and Information Processing* 9 (2020), e15. <https://doi.org/10.1017/ATSIP.2020.13>
- [37] H. Chen, H. YIN, X. Sun, T. Chen, Bogdan G., and K. Musial. 2020. Multi-Level Graph Convolutional Networks for Cross-Platform Anchor Link Prediction. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20)*. Association for Computing Machinery, New York, NY, USA, 1503–1511. <https://doi.org/10.1145/3394486.3403201>
- [38] X. Chen, K. Xiong, Y. Zhang, L. Xia, D. Yin, and J. X. Huang. 2020. Neural Feature-Aware Recommendation with Signed Hypergraph Convolutional Network. *ACM Transactions on Information Systems* 39, 1, Article 8 (2020), 22 pages. <https://doi.org/10.1145/3423322>
- [39] D. Cheng, J. Chen, W. Peng, W. Ye, F. Lv, T. Zhuang, X. Zeng, and X. He. 2022. IHGNN: Interactive Hypergraph Neural Network for Personalized Product Search. In *Proc. of the ACM Web Conference 2022 (WWW '22)*. Association for Computing Machinery, New York, NY, USA, 256–265. <https://doi.org/10.1145/3485447.3511954>
- [40] E. Chien, C. Pan, J. Peng, and O. Milenkovic. 2022. You are AllSet: A Multiset Function Framework for Hypergraph Neural Networks. In *International Conference on Learning Representations (ICLR)*.
- [41] Y. Chu, C. Feng, and C. Guo. 2018. Social-Guided Representation Learning for Images via Deep Heterogeneous Hypergraph Embedding. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*. 1–6. <https://doi.org/10.1109/ICME.2018.8486506>
- [42] Y. Chu, C. Feng, C. Guo, Y. Wang, and J. N. Hwang. 2019. Event2vec: Heterogeneous hypergraph embedding for event data. *IEEE International Conference on Data Mining Workshops (ICDMW)* (2019), 1022–1029. <https://doi.org/10.1109/ICDMW.2018.00147>
- [43] T.-S. Chua, J. Tang, R. Hong, H. Li, Z. Luo, and Y. Zheng. 2009. NUS-WIDE: A Real-World Web Image Database from National University of Singapore. In *Proceedings of the ACM International Conference on Image and Video Retrieval*. Association for Computing Machinery, New York, NY, USA, Article 48, 9 pages. <https://doi.org/10.1145/1646396.1646452>
- [44] F. R. K. Chung. 1997. *Spectral Graph Theory*. American Mathematical Society.
- [45] G. F. de Arruda, G. Petri, and Y. Moreno. 2020. Social contagion models on hypergraphs. *Physical Review Research* 2 (2020), 023032. Issue 2. <https://doi.org/10.1103/PhysRevResearch.2.023032>

- [46] K. Ding, J. Wang, J. Li, D. Li, and H. Liu. 2020. Be More with Less: Hypergraph Attention Networks for Inductive Text Classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 4927–4936. <https://doi.org/10.18653/v1/2020.emnlp-main.399>
- [47] M. Ding, X. Lin, B. Zeng, and Y. Chai. 2021. Hypergraph Neural Networks with Attention Mechanism for Session-based Recommendation. *Journal of Physics: Conference Series* 2082, 1 (2021), 012007. <https://doi.org/10.1088/1742-6596/2082/1/012007>
- [48] Y. Dong, Z. Hu, K. Wang, Y. Sun, and J. Tang. 2020. Heterogeneous Network Representation Learning. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20)*. International Joint Conferences on Artificial Intelligence Organization, 4861–4867. <https://doi.org/10.24963/ijcai.2020/677>
- [49] Y. Dong, W. Sawin, and Y. Bengio. 2020. HNHN: Hypergraph Networks with Hyperedge Neurons. In *Graph Representation Learning and Beyond Workshop at ICML 2020*.
- [50] M. Everingham, . Van Gool, C.K.I. Williams, J. Winn, and A. Zisserman. 2010. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision* 88, 2 (2010), 303–338. <https://doi.org/10.1007/s11263-009-0275-4>
- [51] G. Faggioli, M. Polato, and F. Aiolli. 2020. Recency Aware Collaborative Filtering for Next Basket Recommendation. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '20)*. Association for Computing Machinery, New York, NY, USA, 80–87. <https://doi.org/10.1145/3340631.3394850>
- [52] H. Fan, F. Zhang, Y. Wei, Z. Li, C. Zou, Y. Gao, and Q. Dai. 2021. Heterogeneous Hypergraph Variational Autoencoder for Link Prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 8 (2021), 4125–4138. <https://doi.org/10.1109/TPAMI.2021.3059313>
- [53] B. Fatemi, P. Taslakian, D. Vazquez, and D. Poole. 2021. Knowledge Hypergraphs: Prediction beyond Binary Relations. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI'20)*. Article 303. <https://doi.org/10.5555/3491440.3491743>
- [54] M. R. Felipe-Lucia, A. M. Guerrero, S. M. Alexander, J. Ashander, J. A. Baggio, M. L. Barnes, O. Bodin, A. Bonn, M.-J. Fortin, R. S. Friedman, J. A. Gephart, K. J. Helmstedt, A. A. Keyes, K. Kroetz, F. Massol, M. J.O. Pocock, J. Sayles, R. M. Thompson, S. A. Wood, and L. E. Dee. 2022. Conceptualizing ecosystem services using social–ecological networks. *Trends in Ecology & Evolution* 37, 3 (2022), 211–222. <https://doi.org/10.1016/j.tree.2021.11.012>
- [55] Y. Feng, H. You, Z. Zhang, R. Ji, and Y. Gao. 2019. Hypergraph Neural Networks. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 1 (2019), 3558–3565. <https://doi.org/10.1609/aaai.v33i01.33013558>
- [56] P.W. Frey and D.J. Slate. 1991. Letter recognition using Holland-style adaptive classifiers. *Machine Learning* 6, 2 (1991), 161–182. <https://doi.org/10.1007/BF00114162>
- [57] S. Fu, W. Liu, Y. Zhou, and L. Nie. 2019. HpLapGCN: Hypergraph p-Laplacian graph convolutional networks. *Neurocomputing* 362 (2019), 166–174. <https://doi.org/10.1016/j.neucom.2019.06.068>
- [58] Fukunaga, Yamada, Stone, and Kasai. 1984. A Representation of Hypergraphs in the Euclidean Space. *IEEE Trans. Comput.* C-33, 4 (1984), 364–367. <https://doi.org/10.1109/TC.1984.1676443>
- [59] G. Gallo, G. Longo, S. Pallottino, and S. Nguyen. 1993. Directed hypergraphs and applications. *Discrete Applied Mathematics* 42, 2 (1993), 177–201. [https://doi.org/10.1016/0166-218X\(93\)90045-P](https://doi.org/10.1016/0166-218X(93)90045-P)
- [60] H. Gao and H. Huang. 2018. Deep Attributed Network Embedding. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*. AAAI Press, 3364–3370. <https://doi.org/10.5555/3304222.3304235>
- [61] M. Gao, L. Chen, X. He, and A. Zhou. 2018. BiNE: Bipartite Network Embedding. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (Ann Arbor, MI, USA) (SIGIR '18)*. Association for Computing Machinery, New York, NY, USA, 715–724. <https://doi.org/10.1145/3209978.3209987>
- [62] Y. Gao, Y. Feng, S. Ji, and R. Ji. 2023. HGNN+: General Hypergraph Neural Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45, 3 (2023), 3181–3199. <https://doi.org/10.1109/TPAMI.2022.3182052>
- [63] Y. Gao, Z. Zhang, H. Lin, X. Zhao, S. Du, and C. Zou. 2022. Hypergraph Learning: Methods and Practices. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 5 (2022), 2548–2566. <https://doi.org/10.1109/TPAMI.2020.3039374>
- [64] I. Goodfellow, Y. Bengio, and A. Courville. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- [65] S. Gopalakrishnan, S. Sridharan, S. R. Nayak, J. Nayak, and S. Venkataraman. 2022. Central hubs prediction for bio networks by directed hypergraph - GA with validation to COVID-19 PPI. *Pattern Recognition Letters* 153 (2022), 246–253. <https://doi.org/10.1016/j.patrec.2021.12.015>
- [66] D. Grattarola. 2021. *Graph Neural Networks Operators and Architectures*. Ph. D. Dissertation. Faculty of Informatics of the Università della Svizzera Italiana.
- [67] J. Grilli, G. Barabás, M. J. Michalska-Smith, and S. Allesina. 2017. Higher-order interactions stabilize dynamics in competitive network models. *Nature* 548, 7666 (2017), 210–213. <https://doi.org/10.1038/nature23273>
- [68] A. Grover and J. Leskovec. 2016. Node2vec: Scalable Feature Learning for Networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*. Association for Computing Machinery, New York, NY, USA, 855–864. <https://doi.org/10.1145/2939672.2939754>
- [69] Y. Guan, X. Sun, and Y. Sun. 2021. Sparse relation prediction based on hypergraph neural networks in online social networks. *World Wide Web* (2021). <https://doi.org/10.1007/s11280-021-00936-w>
- [70] Z. Guan, C. Wang, J. Bu, C. Chen, K. Yang, D. Cai, and X. He. 2010. Document Recommendation in Social Tagging Services. In *Proceedings of the 19th International Conference on World Wide Web (WWW '10)*. Association for Computing Machinery, New York, NY, USA, 391–400. <https://doi.org/10.1145/1772690.1772731>

- [71] H. Gui, J. Liu, F. Tao, M. Jiang, B. Norick, and J. Han. 2016. Large-Scale Embedding Learning in Heterogeneous Event Data. *2016 IEEE 16th International Conference on Data Mining (ICDM)* (2016), 907–912. <https://doi.org/10.1109/icdm.2016.0111>
- [72] H. Gui, J. Liu, F. Tao, M. Jiang, B. Norick, L. Kaplan, and J. Han. 2017. Embedding Learning with Events in Heterogeneous Information Networks. *IEEE Transactions on Knowledge and Data Engineering* 29, 11 (2017), 2428–2441. <https://doi.org/10.1109/tkde.2017.2733530>
- [73] L. Guo, H. Yin, T. Chen, X. Zhang, and K. Zheng. 2021. Hierarchical Hyperedge Embedding-Based Representation Learning for Group Recommendation. *ACM Transactions on Information Systems* 40, 1, Article 3 (2021), 27 pages. <https://doi.org/10.1145/3457949>
- [74] X. Guo, L. Zhao, Z. Qin, L. Wu, A. Shehu, and Y. Ye. 2020. Interpretable Deep Graph Generation with Node-Edge Co-Disentanglement. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '20)*. Association for Computing Machinery, New York, NY, USA, 1697–1707. <https://doi.org/10.1145/3394486.3403221>
- [75] W.L. Hamilton. 2020. *Graph Representation Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning, Vol. 14. Morgan & Claypool Publishers. <https://doi.org/10.2200/S01045ED1V01Y202009AIM046>
- [76] W.L. Hamilton, R. Ying, and J. Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 1025–1035. <https://doi.org/10.5555/3294771.3294869>
- [77] W.L. Hamilton, R. Ying, and R. Leskovec. 2017. Representation Learning on Graphs: Methods and Applications. *IEEE Data Engineering Bulletin* 40, 3 (2017), 52–74.
- [78] J. Han, Q. Tao, Y. Tang, and Y. Xia. 2022. DH-HGCN: Dual Homogeneity Hypergraph Convolutional Network for Multiple Social Recommendations. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*. Association for Computing Machinery, New York, NY, USA, 2190–2194. <https://doi.org/10.1145/3477495.3531828>
- [79] F. M. Harper and J. A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Transactions on Interactive Intelligent Systems* 5, 4, Article 19 (2015), 19 pages. <https://doi.org/10.1145/2827872>
- [80] Z.S. Harris. 1954. Distributional Structure. *WORD* 10, 2-3 (1954), 146–162. <https://doi.org/10.1080/00437956.1954.11659520>
- [81] L. He, H. Chen, D. Wang, S. Jameel, P. Yu, and G. Xu. 2021. Click-Through Rate Prediction with Multi-Modal Hypergraphs. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management (CIKM '21)*. Association for Computing Machinery, New York, NY, USA, 690–699. <https://doi.org/10.1145/3459637.3482327>
- [82] D. J. Higham and H.-L. de Kergorlay. 2021. Epidemics on hypergraphs: spectral thresholds for extinction. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 477, 2252 (2021), 20210232. <https://doi.org/10.1098/rspa.2021.0232>
- [83] Z. Hu, J. Wang, S. Chen, and X. Du. 2021. A Semi-Supervised Framework with Efficient Feature Extraction and Network Alignment for User Identity Linkage. In *Proceedings of the 26th International Conference of Database Systems for Advanced Applications (DASFAA 2021)*. Springer-Verlag, Berlin, Heidelberg, 675–691. https://doi.org/10.1007/978-3-030-73197-7_46
- [84] G.-B. Huang, Q.-Y. Zhu, and C.-k. Siew. 2006. Extreme learning machine: Theory and applications. *Neurocomputing* 70, 1 (2006), 489–501. <https://doi.org/10.1016/j.neucom.2005.12.126>
- [85] J. Huang, Ch. Chen, F. Ye, W. Hu, and Z. Zheng. 2020. Nonuniform Hyper-Network Embedding with Dual Mechanism. *ACM Transactions on Information Systems* 38, 3, Article 28 (2020). <https://doi.org/10.1145/3388924>
- [86] J. Huang, C. Chen, F. Ye, J. Wu, Z. Zheng, and G. Ling. 2019. Hyper2vec: Biased Random Walk for Hyper-network Embedding. In *Database Systems for Advanced Applications*, Vol. 11448 LNCS. Springer International Publishing, Cham, 273–277. https://doi.org/10.1007/978-3-030-18590-9_27
- [87] J. Huang, X. Huang, and J. Yang. 2021. Residual Enhanced Multi-Hypergraph Neural Network. In *2021 IEEE International Conference on Image Processing (ICIP)*. 3657–3661. <https://doi.org/10.1109/ICIP42928.2021.9506153>
- [88] J. Huang, F. Lei, S. Wang, S. Wang, and Q. Dai. 2021. Hypergraph Convolutional Network with Hybrid Higher-Order Neighbors. In *Pattern Recognition and Computer Vision*. Springer International Publishing, Cham, 103–114. https://doi.org/10.1007/978-3-030-88013-2_9
- [89] J. Huang, X. Liu, and Y. Song. 2019. Hyper-Path-Based Representation Learning for Hyper-Networks. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM '19)*. Association for Computing Machinery, New York, NY, USA, 449–458. <https://doi.org/10.1145/3357384.3357871>
- [90] J. Huang and J. Yang. 2021. UniGNN: a Unified Framework for Graph and Hypergraph Neural Networks. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21)*. International Joint Conferences on Artificial Intelligence Organization, 2563–2569. <https://doi.org/10.24963/ijcai.2021/353>
- [91] S. Huang, D. Yang, Y. Ge, and X. Zhang. 2016. Discriminant Hyper-Laplacian Projections and its scalable extension for dimensionality reduction. *Neurocomputing* 173 (2016), 145–153. <https://doi.org/10.1016/j.neucom.2015.01.101>
- [92] I. Iacopini, G. Petri, A. Baronchelli, and A. Barrat. 2022. Group interactions modulate critical mass dynamics in social convention. *Communications Physics* 5, 1 (2022), 64. <https://doi.org/10.1038/s42005-022-00845-y>
- [93] B. Jhun, M. Jo, and B. Kahng. 2019. Simplicial SIS model in scale-free uniform hypergraph. *Journal of Statistical Mechanics: Theory and Experiment* 2019, 12 (2019), 123207. <https://doi.org/10.1088/1742-5468/ab5367>
- [94] J. Ji, Y. Ren, and M. Lei. 2022. FC-HAT: Hypergraph Attention Network for Functional Brain Network Classification. *Information Sciences* (2022). <https://doi.org/10.1016/j.ins.2022.07.041>
- [95] S. Ji, Y. Feng, R. Ji, X. Zhao, W. Tang, and Y. Gao. 2020. Dual Channel Hypergraph Collaborative Filtering. In *Proceedings of the International Conference on Knowledge Discovery & Data Mining (KDD '20)*. Association for Computing Machinery, New York, NY, USA, 2020–2029. <https://doi.org/10.1145/3394486.3403253>

- [96] R. Jia, X. Zhou, L. Dong, and S. Pan. 2021. Hypergraph Convolutional Network for Group Recommendation. In *2021 IEEE International Conference on Data Mining (ICDM)*. 260–269. <https://doi.org/10.1109/ICDM51629.2021.00036>
- [97] J. Jiang, Y. Wei, Y. Feng, J. Cao, and Y. Gao. 2019. Dynamic Hypergraph Neural Networks. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-19)*. International Joint Conferences on Artificial Intelligence Organization, 2635–2641. <https://doi.org/10.24963/ijcai.2019/366>
- [98] H. Jin, Y. Wu, H. Huang, Y. Song, H. Wei, and X. Shi. 2022. Modeling Information Diffusion with Sequential Interactive Hypergraphs. *IEEE Transactions on Sustainable Computing* 7, 3 (2022), 644–655. <https://doi.org/10.1109/TSUSC.2022.3152366>
- [99] J. Jo, J. Baek, S. Lee, D. Kim, M. Kang, and S. J. Hwang. 2021. Edge Representation Learning with Hypergraphs. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 7534–7546.
- [100] T. Joachims. 1997. A Probabilistic Analysis of the Rocchio Algorithm with TFIDF for Text Categorization. In *Proceedings of the Fourteenth International Conference on Machine Learning (ICML '97)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 143–151.
- [101] J. Kim, S. Oh, and S. Hong. 2021. Transformers Generalize DeepSets and Can be Extended to Graphs and Hypergraphs. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 28016–28028.
- [102] D. P. Kingma and J. Ba. 2014. Adam: A Method for Stochastic Optimization. <https://doi.org/10.48550/ARXIV.1412.6980>
- [103] T. N. Kipf and M. Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *Proceedings of International Conference on Learning Representations*.
- [104] E.B. Kosmatopoulos, M.M. Polycarpou, M.A. Christodoulou, and P.A. Ioannou. 1995. High-order neural network structures for identification of dynamical systems. *IEEE Transactions on Neural Networks* 6, 2 (1995), 422–431. <https://doi.org/10.1109/72.363477>
- [105] V. La Gatta, V. Moscato, M. Pennone, M. Postiglione, and G. Sperli. 2022. Music Recommendation via Hypergraph Embedding. *IEEE Transactions on Neural Networks and Learning Systems* (2022), 1–13. <https://doi.org/10.1109/TNNLS.2022.3146968>
- [106] N. W. Landry and J. G. Restrepo. 2020. The effect of heterogeneity on hypergraph contagion models. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 30, 10 (2020), 103117. <https://doi.org/10.1063/5.0020034>
- [107] J. Leskovec, L. A. Adamic, and B. A. Huberman. 2007. The Dynamics of Viral Marketing. *ACM Transactions on Web* 1, 1 (2007), 5–es. <https://doi.org/10.1145/1232722.1232727>
- [108] B. Li and D. Pi. 2020. Network representation learning: a systematic literature review. *Neural Computing and Applications* 32, 21 (2020), 16647–16679. <https://doi.org/10.1007/s00521-020-04908-5>
- [109] H. Li, X. Wang, Z. Zhang, Z. Yuan, H. Li, and W. Zhu. 2021. Disentangled Contrastive Learning on Graphs. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 21872–21884.
- [110] Y. Li, H. Chen, X. Sun, Z. Sun, L. Li, L. Cui, P.S. Yu, and G. Xu. 2021. Hyperbolic Hypergraphs for Sequential Recommendation. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management (CIKM '21)*. Association for Computing Machinery, New York, NY, USA, 988–997. <https://doi.org/10.1145/3459637.3482351>
- [111] Y. Li, D. Tarlow, M. Brockschmidt, and R. S. Zemel. 2016. Gated Graph Sequence Neural Networks. In *International Conference on Learning Representations (ICLR)*.
- [112] X. Liao, Y. Xu, and H. Ling. 2021. Hypergraph Neural Networks for Hypergraph Matching. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. 1246–1255. <https://doi.org/10.1109/ICCV48922.2021.00130>
- [113] D. Liben-Nowell and J. Kleinberg. 2007. The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology* 58, 7 (2007), 1019–1031. <https://doi.org/10.1002/asi.20591>
- [114] B. Liu, P. Zhao, F. Zhuang, X. Xian, Y. Liu, and V. S. Sheng. 2021. Knowledge-Aware Hypergraph Neural Network for Recommender Systems. In *Database Systems for Advanced Applications*. Springer International Publishing, Cham, 132–147. https://doi.org/10.1007/978-3-030-73200-4_9
- [115] Z. Liu, Z. Zhang, Y. Cai, Y. Miao, and Z. Chen. 2021. Semi-Supervised Classification via Hypergraph Convolutional Extreme Learning Machine. *Applied Sciences* 11, 9 (2021). <https://doi.org/10.3390/app11093867>
- [116] Q. F. Lotito, M. Contisciani, C. De Bacco, L. Di Gaetano, L. Gallo, A. Montresor, F. Musciotto, N. Ruggeri, and F. Battiston. 2023. Hypergraphx: a library for higher-order network analysis. *Journal of Complex Networks* 11, 3 (2023). <https://doi.org/10.1093/comnet/cnad019>
- [117] Q. F. Lotito, F. Musciotto, A. Montresor, and F. Battiston. 2022. Higher-order motif analysis in hypergraphs. *Communications Physics* 5, 1 (2022), 79. <https://doi.org/10.1038/s42005-022-00858-7>
- [118] F. Luo, B. Du, L. Zhang, L. Zhang, and D. Tao. 2019. Feature Learning Using Spatial-Spectral Hypergraph Discriminant Analysis for Hyperspectral Image. *IEEE Trans. on Cybernetics* 49, 7 (2019), 2406–2419. <https://doi.org/10.1109/TCYB.2018.2810806>
- [119] F. Luo, G. Guo, Z. Lin, J. Ren, and X. Zhou. 2020. Semisupervised Hypergraph Discriminant Learning for Dimensionality Reduction of Hyperspectral Image. *IEEE J. of Selected Topics in Applied Earth Observations and Remote Sensing* 13 (2020), 4242–4256. <https://doi.org/10.1109/JSTARS.2020.3011431>
- [120] X. Luo, J. Peng, and J. Liang. 2022. Directed hypergraph attention network for traffic forecasting. *IET Intelligent Transport Systems* 16, 1 (2022), 85–98. <https://doi.org/10.1049/itr2.12130>
- [121] Jichao M., Yanjiang W., Baodi L., and Weifeng L. 2021. Accurately modeling the human brain functional correlations with hypergraph Laplacian. *Neurocomputing* 428 (2021), 239–247. <https://doi.org/10.1016/j.neucom.2020.11.021>
- [122] T. Ma and J. Guo. 2018. Study on information transmission model of enterprise informal organizations based on the hypernetwork. *Chinese Journal of Physics* 56, 5 (2018), 2424 – 2438. <https://doi.org/10.1016/j.cjph.2018.06.018>

- [123] X. Ma, W. Liu, Q. Tian, and Y. Gao. 2022. Learning Representation on Optimized High-Order Manifold for Visual Classification. *IEEE Transactions on Multimedia* 24 (2022), 3989–4001. <https://doi.org/10.1109/TMM.2021.3111500>
- [124] Z. Ma, Z. Jiang, and H. Zhang. 2021. Hyperspectral image classification using spectral-spatial hypergraph convolution neural network. In *Image and Signal Processing for Remote Sensing XXVII*, Vol. 11862. SPIE, 118620I. <https://doi.org/10.1117/12.2599787>
- [125] Z. Ma, Z. Jiang, and H. Zhang. 2022. Hyperspectral Image Classification Using Feature Fusion Hypergraph Convolution Neural Network. *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022), 1–14. <https://doi.org/10.1109/TGRS.2021.3123423>
- [126] M. Madine, I. Rezik, and N. Werghi. 2020. Diagnosing Autism Using T1-W MRI With Multi-Kernel Learning and Hypergraph Neural Network. In *2020 IEEE International Conference on Image Processing (ICIP)*. 438–442. <https://doi.org/10.1109/ICIP40778.2020.9190924>
- [127] S. Maleki, D.P. Wall, and K. Pingali. 2021. NetVec: A Scalable Hypergraph Embedding System. In *Proceedings of the International Conference on Machine Learning - Workshops*.
- [128] V. Martínez, F. Berzal, and J.-C. Cubero. 2016. A Survey of Link Prediction in Complex Networks. *Comput. Surveys* 49, 4, Article 69 (2016), 33 pages. <https://doi.org/10.1145/3012704>
- [129] T. Mikolov, K. Chen, G. Corrado, and J. Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *Proceedings of International Conference on Learning Representation*.
- [130] George A. Miller. 1994. WordNet: A Lexical Database for English. In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*. <https://aclanthology.org/H94-1111>
- [131] L. Neuhäuser, R. Lambiotte, and M.T. Schaub. 2021. Consensus dynamics on temporal hypergraphs. *Physics Review E* 104 (2021), 064305. Issue 6. <https://doi.org/10.1103/PhysRevE.104.064305>
- [132] L. Neuhäuser, A. Mellor, and R. Lambiotte. 2020. Multibody interactions and nonlinear consensus dynamics on networked systems. *Physics Review E* 101 (2020), 032310. Issue 3. <https://doi.org/10.1103/PhysRevE.101.032310>
- [133] J. Ni, J. Li, and J. McAuley. 2019. Justifying Recommendations using Distantly-Labeled Reviews and Fine-Grained Aspects. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 188–197. <https://doi.org/10.18653/v1/D19-1018>
- [134] L. Nong, J. Wang, J. Lin, H. Qiu, L. Zheng, and W. Zhang. 2021. Hypergraph Wavelet Neural Networks for 3D Object Classification. *Neurocomputing* 463, C (2021), 580–595. <https://doi.org/10.1016/j.neucom.2021.08.006>
- [135] M. Ou, P. Cui, J. Pei, Z. Zhang, and W. Zhu. 2016. Asymmetric Transitivity Preserving Graph Embedding. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*. Association for Computing Machinery, New York, NY, USA, 1105–1114. <https://doi.org/10.1145/2939672.2939751>
- [136] S. Pang, K. Zhang, S. Wang, Y. Zhang, S. He, W. Wu, and S. Qiao. 2021. HGDD: A Drug-Disease High-Order Association Information Extraction Method for Drug Repurposing via Hypergraph. In *Bioinformatics Research and Applications*. Springer International Publishing, Cham, 424–435. https://doi.org/10.1007/978-3-030-91415-8_36
- [137] R. M. Pattanayak and H. S. Behera. 2018. Higher Order Neural Network and Its Applications: A Comprehensive Survey. In *Progress in Computing, Analytics and Networking*. Springer Singapore, Singapore, 695–709. https://doi.org/10.1007/978-981-10-7871-2_67
- [138] J. Payne. 2019. Deep Hyperedges: a Framework for Transductive and Inductive Learning on Hypergraphs. In *Proceedings of Neural Information Processing Systems*.
- [139] D. Peng and S. Zhang. 2022. GC-HGNN: A global-context supported hypergraph neural network for enhancing session-based recommendation. *Electronic Commerce Research and Applications* 52 (2022), 101129. <https://doi.org/10.1016/j.elerap.2022.101129>
- [140] B. Perozzi, R. Al-Rfou, and S. Skiena. 2014. DeepWalk: Online Learning of Social Representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '14)*. Association for Computing Machinery, New York, NY, USA, 701–710. <https://doi.org/10.1145/2623330.2623732>
- [141] L. Pu and B. Faltings. 2012. Hypergraph Learning with Hyperedge Expansion. In *Machine Learning and Knowledge Discovery in Databases*. Springer Berlin Heidelberg, Berlin, Heidelberg, 410–425. https://doi.org/10.1007/978-3-642-33460-3_32
- [142] P. Ren, R. C. Wilson, and E. R. Hancock. 2008. Spectral Embedding of Feature Hypergraphs. In *Structural, Syntactic, and Statistical Pattern Recognition*. 308–317. https://doi.org/10.1007/978-3-540-89689-0_35
- [143] S. Rendle. 2010. Factorization Machines. In *Proceedings of the International Conference on Data Mining*. 995–1000. <https://doi.org/10.1109/ICDM.2010.127>
- [144] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI '09)*. AUAI Press, Arlington, Virginia, USA, 452–461.
- [145] F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor. 2010. *Recommender Systems Handbook* (1st ed.). Springer-Verlag, Berlin, Heidelberg.
- [146] J.A. Rodriguez. 2002. On the Laplacian Eigenvalues and Metric Parameters of Hypergraphs. *Linear and Multilinear Algebra* 50, 1 (2002), 1–14. <https://doi.org/10.1080/03081080290011692>
- [147] R. Rossi and N. Ahmed. 2015. The Network Data Repository with Interactive Graph Analytics and Visualization. In *AAAI Conference on Artificial Intelligence*, Vol. 29. AAAI, New York, NY, USA, 4292–4293.
- [148] P. Rosso, D. Yang, and P. Cudré-Mauroux. 2020. Beyond Triplets: Hyper-Relational Knowledge Graph Embedding for Link Prediction. In *Proceedings of The Web Conference 2020 (WWW '20)*. Association for Computing Machinery, New York, NY, USA, 1885–1896. <https://doi.org/10.1145/3366423.3380257>

- [149] R. Sahasrabudde, L. Neuhäuser, and R. Lambiotte. 2021. Modelling non-linear consensus dynamics on hypergraphs. *Journal of Physics: Complexity* 2, 2 (feb 2021), 025006. <https://doi.org/10.1088/2632-072X/abcea3>
- [150] S. Saito, D. P. Mandic, and H. Suzuki. 2018. Hypergraph p-Laplacian: A Differential Geometry View. *Proceedings of the AAAI Conference on Artificial Intelligence* 32, 1 (2018). <https://doi.org/10.1609/aaai.v32i1.11823>
- [151] A. Sanchez-Gorostiaga, D. Bajić, M. L. Osborne, J. F. Poyatos, and A. Sanchez. 2019. High-order interactions distort the functional landscape of microbial consortia. *PLOS Biology* 17, 12 (12 2019), 1–34. <https://doi.org/10.1371/journal.pbio.3000550>
- [152] N. Sasikaladevi. 2022. Robust and fast Plant Pathology Prognostics (P3) tool based on deep convolutional neural network. *Multimedia Tools and Applications* 81, 5 (2022), 7271–7283. <https://doi.org/10.1007/s11042-022-11902-7>
- [153] R. Sawhney, S. Agarwal, A. Wadhwa, T. Derr, and R. R. Shah. 2021. Stock Selection via Spatiotemporal Hypergraph Attention Network: A Learning to Rank Approach. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 1 (2021), 497–504. <https://doi.org/10.1609/aaai.v35i1.16127>
- [154] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, and T. Eliassi-Rad. 2008. Collective Classification in Network Data. *AI Magazine* 29, 3 (2008), 93. <https://doi.org/10.1609/aimag.v29i3.2157>
- [155] J. Shun. 2020. Practical Parallel Hypergraph Algorithms. In *Proceedings of the ACM Symposium on Principles and Practice of Parallel Programming*, 232–249. <https://doi.org/10.1145/3332466.3374527>
- [156] B. Srinivasan, D. Zheng, and G. Karypis. 2021. Learning over Families of Sets - Hypergraph Representation Learning for Higher Order Tasks. In *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*. Siam Society, 756–764. <https://doi.org/10.1137/1.9781611976700.85>
- [157] W.N. Street, W.H. Wolberg, and O.L. Mangasarian. 1993. Nuclear feature extraction for breast tumor diagnosis. In *Biomedical Image Processing and Biomedical Visualization*, Raj S. Acharya and Dmitry B. Goldgof (Eds.), Vol. 1905. International Society for Optics and Photonics, SPIE, 861 – 870. <https://doi.org/10.1117/12.148698>
- [158] L. Sun, S. Ji, and J. Ye. 2008. Hypergraph Spectral Learning for Multi-Label Classification. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining (KDD '08)*. Association for Computing Machinery, New York, NY, USA, 668–676. <https://doi.org/10.1145/1401890.1401971>
- [159] X. Sun, H. Yin, B. Liu, H. Chen, J. Cao, Y. Shao, and N.Q. Viet Hung. 2021. Heterogeneous Hypergraph Embedding for Graph Classification. In *Proceedings of the International Conference on Web Search and Data Mining (WSDM '21)*. Association for Computing Machinery, New York, NY, USA, 725–733. <https://doi.org/10.1145/3437963.3441835>
- [160] X. Sun, H. Yin, B. Liu, H. Chen, Q. Meng, W. Han, and J. Cao. 2021. Multi-Level Hyperedge Distillation for Social Linking Prediction on Sparsely Observed Networks. In *Proc. of the Web Conference 2021 (WWW '21)*. Association for Computing Machinery, New York, NY, USA, 2934–2945. <https://doi.org/10.1145/3442381.3449912>
- [161] X. Sun, H. Yin, B. Liu, Q. Meng, J. Cao, A. Zhou, and H. Chen. 2022. Structure Learning Via Meta-Hyperedge for Dynamic Rumor Detection. *IEEE Transactions on Knowledge and Data Engineering* (2022), 1–12. <https://doi.org/10.1109/TKDE.2022.3221438>
- [162] Y. Sun, Sujuan Wang, Qingshan Liu, Renlong Hang, and Guangcan Liu. 2017. Hypergraph Embedding for Spatial-Spectral Joint Feature Extraction in Hyperspectral Images. *Remote Sensing* 9, 5 (2017). <https://doi.org/10.3390/rs9050506>
- [163] Q. Suo, J. Guo, and A. Shen. 2018. Information spreading dynamics in hypernetworks. *Physica A: Statistical Mechanics and its Applications* 495 (2018), 475 – 487. <https://doi.org/10.1016/j.physa.2017.12.108>
- [164] J. Sybrandt and I. Safro. 2019. FOBE and HOBE: First- and High-Order Bipartite Embeddings. arXiv, abs/1905.10953.
- [165] J. Sybrandt, R. Shaydulin, and I. Safro. 2020. Hypergraph Partitioning With Embeddings. *IEEE Transactions on Knowledge and Data Engineering* 34, 6 (2020), 2771–2782. <https://doi.org/10.1109/TKDE.2020.3017120>
- [166] T. Thonet, J.-M. Renders, M. Choi, and J. Kim. 2022. Joint Personalized Search and Recommendation with Hypergraph Convolutional Networks. In *Advances in Information Retrieval (ECIR 2022)*. Springer-Verlag, Berlin, Heidelberg, 443–456. https://doi.org/10.1007/978-3-030-99736-6_30
- [167] H.T. Trung, T. Van Vinh, N.T. Tam, J. Jo, H. Yin, and N.Q.V. Hung. 2022. Learning Holistic Interactions in LBSNs with High-order, Dynamic, and Multi-role Contexts. *IEEE Transactions on Knowledge and Data Engineering* 35, 5 (2022), 5002–5016. <https://doi.org/10.1109/TKDE.2022.3150792>
- [168] K. Tu, P. Cui, X. Wang, F. Wang, and W. Zhu. 2018. Structural Deep Embedding for Hyper-Networks. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI'18)*. AAAI Press, Article 53. <https://doi.org/10.5555/3504035.3504088>
- [169] L. van der Maaten and G. Hinton. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research* 9 (2008), 2579–2605.
- [170] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, Vol. 30. Curran Associates, Inc.
- [171] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. 2017. Graph Attention Networks. *6th International Conference on Learning Representations* (2017).
- [172] M. Vijaikumar, D. Hada, and S. Shevade. 2021. HyperTeNet: Hypergraph and Transformer-based Neural Network for Personalized List Continuation. In *2021 IEEE International Conference on Data Mining (ICDM)*. 1210–1215. <https://doi.org/10.1109/ICDM51629.2021.00146>
- [173] H. Wang, B. Chen, and W.-J. Li. 2013. Collaborative Topic Regression with Social Regularization for Tag Recommendation. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (Beijing, China) (IJCAI '13)*. AAAI Press, 2719–2725. <https://doi.org/10.5555/2540128.2540520>
- [174] J. Wang and J. Caverlee. 2019. Recurrent Recommendation with Local Coherence. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining (Melbourne VIC, Australia) (WSDM '19)*. Association for Computing Machinery, New York, NY, USA, 564–572. <https://doi.org/10.1145/3289600.3291024>

- [175] J. Wang, K. Ding, L. Hong, H. Liu, and J. Caverlee. 2020. Next-Item Recommendation with Sequential Hypergraphs. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*. ACM, New York, NY, USA, 1101–1110. <https://doi.org/10.1145/3397271.3401133>
- [176] J. Wang, K. Ding, Z. Zhu, and J. Caverlee. 2021. Session-based Recommendation with Hypergraph Attention Networks. In *Proc. of the SIAM International Conference on Data Mining*. Society for Industrial and Applied Mathematics, 82–90. <https://doi.org/10.1137/1.9781611976700.10>
- [177] J. Wang, Y. Zhang, L. Wang, Y. Hu, X. Piao, and B. Yin. 2022. Multitask Hypergraph Convolutional Networks: A Heterogeneous Traffic Prediction Framework. *IEEE Transactions on Intelligent Transportation Systems* (2022), 1–11. <https://doi.org/10.1109/TITS.2022.3168879>
- [178] J. Wang, Y. Zhang, Y. Wei, Y. Hu, X. Piao, and B. Yin. 2021. Metro Passenger Flow Prediction via Dynamic Hypergraph Convolution Networks. *IEEE Transactions on Intelligent Transportation Systems* 22, 12 (2021), 7891–7903. <https://doi.org/10.1109/TITS.2021.3072743>
- [179] N. Wang, S. Wang, Y. Wang, Q. Z. Sheng, and M. A. Orgun. 2022. Exploiting intra- and inter-session dependencies for session-based recommendations. *World Wide Web* 25, 1 (2022), 425–443. <https://doi.org/10.1007/s11280-021-00930-2>
- [180] R. Wang, Y. Li, S. Lin, H. Xie, Y. Xu, and J. C. S. Lui. 2021. On Modeling Influence Maximization in Social Activity Networks under General Settings. *ACM Transactions on Knowledge Discovery from Data* 15, 6, Article 108 (2021), 28 pages. <https://doi.org/10.1145/3451218>
- [181] X. Wang, D. Bo, C. Shi, S. Fan, Y. Ye, and P. S. Yu. 2022. A Survey on Heterogeneous Graph Embedding: Methods, Techniques, Applications and Sources. *IEEE Transactions on Big Data* (2022), 1–1. <https://doi.org/10.1109/TBDDATA.2022.3177455>
- [182] L. Wu, D. Wang, K. Song, S. Feng, Y. Zhang, and G. Yu. 2021. Dual-view hypergraph neural networks for attributed graph learning. *Knowledge-Based Systems* 227 (2021), 107185. <https://doi.org/10.1016/j.knsys.2021.107185>
- [183] X. Wu, Q. Chen, W. Li, Y. Xiao, and B. Hu. 2020. AdaHGNN: Adaptive Hypergraph Neural Networks for Multi-Label Image Classification. In *Proceedings of the International Conference on Multimedia (MM '20)*. ACM, New York, NY, USA, 284–293. <https://doi.org/10.1145/3394171.3414046>
- [184] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. 2021. A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems* 32, 1 (2021), 4–24. <https://doi.org/10.1109/TNNLS.2020.2978386>
- [185] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 2015. 3D ShapeNets: A Deep Representation for Volumetric Shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [186] L. Xia, C. Huang, Y. Xu, J. Zhao, D. Yin, and J. Huang. 2022. Hypergraph Contrastive Collaborative Filtering. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (Madrid, Spain) (SIGIR '22)*. Association for Computing Machinery, New York, NY, USA, 70–79. <https://doi.org/10.1145/3477495.3532058>
- [187] L. Xia, C. Huang, and C. Zhang. 2022. Self-Supervised Hypergraph Transformer for Recommender Systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22)*. Association for Computing Machinery, New York, NY, USA, 2100–2109. <https://doi.org/10.1145/3534678.3539473>
- [188] L. Xia, P. Zheng, X. Huang, and C. Liu. 2021. A novel hypergraph convolution network-based approach for predicting the material removal rate in chemical mechanical planarization. *Journal of Intelligent Manufacturing* 33, 8 (2021), 2295–2306. <https://doi.org/10.1007/s10845-021-01784-1>
- [189] L. Xia, P. Zheng, and C. Liu. 2021. Predicting the material removal rate in chemical mechanical planarization process: A hypergraph neural network-based approach. *Proceedings of the ASME Design Engineering Technical Conference 2* (2021). <https://doi.org/10.1115/DETC2021-68250>
- [190] X. Xia, H. Yin, J. Yu, Q. Wang, L. Cui, and X. Zhang. 2021. Self-Supervised Hypergraph Convolutional Networks for Session-based Recommendation. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 5 (2021), 4503–4511. <https://doi.org/10.1609/aaai.v35i5.16578>
- [191] B. Xu, N. Wang, T. Chen, and M. Li. 2015. Empirical Evaluation of Rectified Activations in Convolutional Network. <https://doi.org/10.48550/ARXIV.1505.00853>
- [192] G. Xue, M. Zhong, J. Li, J. Chen, C. Zhai, and R. Kong. 2022. Dynamic network embedding survey. *Neurocomputing* 472 (2022), 212–223. <https://doi.org/10.1016/j.neucom.2021.03.138>
- [193] H. Xue, L. Yang, V. Rajan, W. Jiang, Y. Wei, and Y. Lin. 2021. Multiplex Bipartite Network Embedding Using Dual Hypergraph Convolutional Networks. In *Proceedings of the Web Conference 2021 (WWW '21)*. ACM, New York, NY, USA, 1649–1660. <https://doi.org/10.1145/3442381.3449954>
- [194] N. Yadati. 2020. Neural Message Passing for Multi-Relational Ordered and Recursive Hypergraphs. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NIPS'20)*. Curran Associates Inc., Red Hook, NY, USA, Article 276.
- [195] N. Yadati, M. Nimishakavi, P. Yadav, V. Nitin, A. Louis, and P. Talukdar. 2019. HyperGCN: A New Method For Training Graph Convolutional Networks on Hypergraphs. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc.
- [196] N. Yadati, V. Nitin, M. Nimishakavi, P. Yadav, A. Louis, and P. Talukdar. 2020. NHP: Neural Hypergraph Link Prediction. In *Proceedings of the International Conference on Information & Knowledge Management (CIKM '20)*. Association for Computing Machinery, New York, NY, USA, 1705–1714. <https://doi.org/10.1145/3340531.3411870>
- [197] S. Yan, Z. Zhang, X. Sun, G. Xu, L. Jin, and S. Li. 2022. HYPER2: Hyperbolic embedding for hyper-relational link prediction. *Neurocomputing* 492 (2022), 440–451. <https://doi.org/10.1016/j.neucom.2022.04.026>
- [198] C. Yang, R. Wang, S. Yao, and T. Abdelzaher. 2022. Semi-Supervised Hypergraph Node Classification on Hypergraph Line Expansion. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management (CIKM '22)*. Association for Computing Machinery, New York, NY, USA, 2352–2361. <https://doi.org/10.1145/3511808.3557447>
- [199] D. Yang, B. Qu, J. Yang, and P. Cudre-Mauroux. 2019. Revisiting User Mobility and Social Relationships in LBSNs: A Hypergraph Embedding Approach. In *The World Wide Web Conference (WWW '19)*. Association for Computing Machinery, New York, NY, USA, 2147–2157. <https://doi.org/10.1145/3308558.3313635>

- [200] D. Yang, B. Qu, J. Yang, and P. Cudre-Mauroux. 2020. LBSN2Vec++: Heterogeneous Hypergraph Embedding for Location-Based Social Networks. *IEEE Transactions on Knowledge and Data Engineering* 34 (2020), 1843–1855. <https://doi.org/10.1109/tkde.2020.2997869>
- [201] J. Yang and J. Leskovec. 2015. Defining and Evaluating Network Communities Based on Ground-Truth. *Knowledge Information Systems* 42, 1 (2015), 181–213. <https://doi.org/10.1007/s10115-013-0693-z>
- [202] J. Yi and J. Park. 2020. Hypergraph Convolutional Recurrent Neural Network. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '20)*. Association for Computing Machinery, New York, NY, USA, 3366–3376. <https://doi.org/10.1145/3394486.3403389>
- [203] S. Yoon, H. Song, K. Shin, and Y. Yi. 2020. How Much and When Do We Need Higher-Order Information in Hypergraphs? A Case Study on Hyperedge Prediction. In *Proceedings of The Web Conference 2020 (WWW '20)*. ACM, New York, NY, USA, 2627–2633. <https://doi.org/10.1145/3366423.3380016>
- [204] C. Yu, C. Tai, T. Chan, and Y. Yang. 2018. Modeling Multi-way Relations with Hypergraph Embedding. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM'18)*. ACM, New York, NY, USA, 1707–1710. <https://doi.org/10.1145/3269206.3269274>
- [205] Guihai Yu, Xiying Yuan, and Hui Qu. 2019. Signed k-uniform hypergraphs and tensors. *Linear Algebra Appl.* 580 (2019), 1–13. <https://doi.org/10.1016/j.laa.2019.06.011>
- [206] J. Yu, H. Yin, J. Li, Q. Wang, N.Q.V. Hung, and X. Zhang. 2021. Self-Supervised Multi-Channel Hypergraph Convolutional Network for Social Recommendation. In *Proc. of the Web Conference 2021 (WWW '21)*. ACM, New York, NY, USA, 413–424. <https://doi.org/10.1145/3442381.3449844>
- [207] H. Yuan and Y. Y. Tang. 2015. Learning With Hypergraph for Hyperspectral Image Feature Extraction. *IEEE Geoscience and Remote Sensing Letters* 12, 8 (2015), 1695–1699. <https://doi.org/10.1109/LGRS.2015.2419713>
- [208] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Poczos, R.R. Salakhutdinov, and A.J. Smola. 2017. Deep Sets. In *Advances in Neural Information Processing Systems* 30. Curran Associates, Inc., 3391–3401.
- [209] E. Zangerle, M. Pichl, W. Gassler, and G. Specht. 2014. #nowplaying Music Dataset: Extracting Listening Behavior from Twitter. In *Proceedings of the First International Workshop on Internet-Scale Multimedia Management (WISMM '14)*. Association for Computing Machinery, New York, NY, USA, 21–26. <https://doi.org/10.1145/2661714.2661719>
- [210] D. Zhang, J. Yin, X. Zhu, and C. Zhang. 2020. Network Representation Learning: A Survey. *IEEE Transactions on Big Data* 6, 1 (2020), 3–28. <https://doi.org/10.1109/TBDATA.2018.2850013>
- [211] J. Zhang, M. Gao, J. Yu, L. Guo, J. Li, and H. Yin. 2021. Double-Scale Self-Supervised Hypergraph Learning for Group Recommendation. In *Proceedings of the Int. Conf. on Information & Knowledge Management (CIKM '21)*. Association for Computing Machinery, New York, NY, USA, 2557–2567. <https://doi.org/10.1145/3459637.3482426>
- [212] L. Zhang, J. Guo, J. Wang, J. Wang, S. Li, and C. Zhang. 2022. Hypergraph and Uncertain Hypergraph Representation Learning Theory and Methods. *Mathematics* 10, 11 (2022). <https://doi.org/10.3390/math10111921>
- [213] M. Zhang, H. Luo, W. Song, H. Mei, and C. Su. 2021. Spectral-Spatial Offset Graph Convolutional Networks for Hyperspectral Image Classification. *Remote Sensing* 13, 21 (2021). <https://doi.org/10.3390/rs13214342>
- [214] R. Zhang and J. Ma. 2020. MATCHA: Probing Multi-way Chromatin Interaction with Hypergraph Representation Learning. *Cell Systems* 10, 5 (2020), 397–407.e5. <https://doi.org/10.1016/j.cels.2020.04.004>
- [215] R. Zhang and J. Ma. 2020. Probing Multi-way Chromatin Interaction with Hypergraph Representation Learning. In *Research in Computational Molecular Biology*. Springer International Publishing, Cham, 276–277. https://doi.org/10.1007/978-3-030-45257-5_37
- [216] R. Zhang, T. Zhou, and J. Ma. 2022. Multiscale and integrative single-cell Hi-C analysis with Higashi. *Nature Biotechnology* 40, 2 (2022), 254–261. <https://doi.org/10.1038/s41587-021-01034-y>
- [217] R. Zhang, Y. Zou, and J. Ma. 2020. Hyper-SAGNN: a self-attention based graph neural network for hypergraphs. In *International Conference on Learning Representations (ICLR)*.
- [218] S. Zhang, H. Tong, J. Xu, and R. Maciejewski. 2019. Graph convolutional networks: a comprehensive review. *Computational Social Networks* 6, 1 (2019), 11. <https://doi.org/10.1186/s40649-019-0069-y>
- [219] Q. Zheng and D. B. Skillicorn. 2015. Spectral Embedding of Directed Networks. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*. Association for Computing Machinery, New York, NY, USA, 432–439. <https://doi.org/10.1145/2808797.2809310>
- [220] V.W. Zheng, B. Cao, Y. Zheng, X. Xie, and Q. Yang. 2010. Collaborative Filtering Meets Mobile Recommendation: A User-Centered Approach. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI'10)*. AAAI Press, 236–241. <https://doi.org/10.5555/2898607.2898645>
- [221] D. Zhou, J. Huang, and B. Schölkopf. 2007. Learning with Hypergraphs: Clustering, Classification, and Embedding. In *Proceedings of Neural Information Processing Systems*. 1601–1608.
- [222] J. Zhu, S. Ghosh, and W. Wu. 2019. Group Influence Maximization Problem in Social Networks. *IEEE Transactions on Computational Social Systems* 6, 6 (2019), 1156–1164. <https://doi.org/10.1109/TCSS.2019.2938575>
- [223] J. Zhu, J. Zhu, J. Ghosh, W. Wu, and J. Yuan. 2019. Social Influence Maximization in Hypergraph in Social Networks. *IEEE Transactions on Network Science and Engineering* 6, 4 (2019), 801–811. <https://doi.org/10.1109/TNSE.2018.2873759>
- [224] Y. Zhu, Z. Guan, T. Tan, H. Liu, D. Cai, and X. He. 2016. Heterogeneous hypergraph embedding for document recommendation. *Neurocomputing* 216 (2016), 150–162. <https://doi.org/10.1016/j.neucom.2016.07.030>
- [225] Y. Zhu and H. Zhao. 2022. Hypernetwork Representation Learning with the Set Constraint. *Applied Sciences* 12, 5 (2022). <https://doi.org/10.3390/app12052650>