Computing Graph Neural Networks: A Survey from Algorithms to Accelerators

SERGI ABADAL, AKSHAY JAIN, ROBERT GUIRADO, JORGE LÓPEZ-ALONSO, and EDUARD ALARCÓN, Universitat Politècnica de Catalunya, Spain

Graph Neural Networks (GNNs) have exploded onto the machine learning scene in recent years owing to their capability to model and learn from graph-structured data. Such an ability has strong implications in a wide variety of fields whose data are inherently relational, for which conventional neural networks do not perform well. Indeed, as recent reviews can attest, research in the area of GNNs has grown rapidly and has led to the development of a variety of GNN algorithm variants as well as to the exploration of ground-breaking applications in chemistry, neurology, electronics, or communication networks, among others. At the current stage research, however, the efficient processing of GNNs is still an open challenge for several reasons. Besides of their novelty, GNNs are hard to compute due to their dependence on the input graph, their combination of dense and very sparse operations, or the need to scale to huge graphs in some applications. In this context, this article aims to make two main contributions. On the one hand, a review of the field of GNNs is presented from the perspective of computing. This includes a brief tutorial on the GNN fundamentals, an overview of the evolution of the field in the last decade, and a summary of operations carried out in the multiple phases of different GNN algorithm variants. On the other hand, an in-depth analysis of current software and hardware acceleration schemes is provided, from which a hardware-software, graph-aware, and communication-centric vision for GNN accelerators is distilled.

CCS Concepts: • **Computing methodologies** → **Machine learning algorithms**; • **Computer systems organization** → **Neural networks**; *Data flow architectures*; • **Mathematics of computing** → **Graph algorithms**; • **Hardware** → *Hardware accelerators*;

Additional Key Words and Phrases: Graph neural networks, GNN algorithms, accelerators, graph embeddings

ACM Reference format:

Sergi Abadal, Akshay Jain, Robert Guirado, Jorge López-Alonso, and Eduard Alarcón. 2021. Computing Graph Neural Networks: A Survey from Algorithms to Accelerators. *ACM Comput. Surv.* 54, 9, Article 191 (October 2021), 38 pages.

<https://doi.org/10.1145/3477141>

1 INTRODUCTION

Machine Learning (ML) has taken the world by storm and has become a fundamental pillar of engineering due to its capacity to solve extremely complex problems, to detect intricate

This work is possible thanks to funding from the European Union's Horizon 2020 research and innovation programme under Grant No. 863337 (WiPLASH project) and the Spanish Ministry of Economy and Competitiveness under contract TEC2017-90034-C2-1-R (ALLIANCE project) that receives funding from FEDER.

Authors' addresses: S. Abadal, A. Jain, R. Guirado, J. López-Alonso, and E. Alarcón, Universitat Politècnica de Catalunya, Spain; emails: abadal@ac.upc.edu, akshay.jain@upc.edu, roberto.guirado.linan@upc.edu, jorlopezalonso@gmail.com, eduard.alarcon@upc.edu.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Copyright held by the owner/author(s).

0360-0300/2021/10-ART191

<https://doi.org/10.1145/3477141>

Fig. 1. GNNs as enablers of a plethora of applications in fields that hinge on graph-structured data.

features in oceans of data, or to automatically generate alternatives that outperform well-engineered, well-known, carefully optimized solutions. As a result, the last decade has witnessed an explosive growth in the use of **Deep Neural Networks (DNNs)** in pursuit of exploiting the advantages of ML in virtually every aspect of our lives [92]: computer vision [67], natural language processing [171], medicine [43], or economics [62] are just a few examples.

However, and in spite of its all-pervasive applicability and potential, it is well known that not all neural network architectures fit to all problems [11]. DNNs take the input data and attempt to extract knowledge taking into account the inductive bias that the connection architecture of the DNN generates. This, in essence, means that the number of DNN layers and their pre-assumed connections determines its suitability to certain tasks. For instance, by not making any assumption on the structure of the data, conventional fully connected neural networks are able to master a wide range of tasks at the cost of being less efficient in general than other DNNs [14]. In contrast, techniques such as **Convolutional Neural Networks (CNNs)** or **Recursive Neural Networks (RNNs)** are biased toward extracting knowledge from the locality and temporal sequentiality of data. This makes them a better fit for specific tasks such as image recognition or treatment of temporal signals, yet incapable of efficiently handling data with arbitrary structures [149].

In light of the above, there has been a recent interest in deep learning techniques able to model graph-structured data [2, 11, 16, 49, 54, 181]. This structure is inherent to a plethora of problems in the field of complex systems in general, and applicable to particular fields such as communication networks where the topology and routing decisions determine its performance [126], synthetic chemistry where molecular structures determine the compound properties [56], social networks where emergent behavior can arise through personal relations [116], or neuroscience where specific connections between neuron types and brain areas determine brain function [100], among many others.

Graph Neural Networks (GNNs) are a set of connectivity-driven models that, since the late 2000s, have been addressing the need for geometric deep learning [57, 130]. In essence, GNN adapt their structure to that of an input graph and, through an iterative process of aggregation of information across vertices, capture the complex dependencies of the underlying system. This allows to predict properties for specific nodes, connections, or the graph as a whole, and generalize to unseen graphs. Due to these powerful features, many relevant applications such as molecule property prediction [47], recommender systems [44], natural language processing [171], traffic speed prediction [161], critical data classification [170], computer vision [152], particle physics [80], and resource allocation in computer networks [125] already utilize GNNs to accomplish their tasks.

For all these reasons, recent years have seen a rapid increase in research activity in the field of GNNs (see Figure 6). Intense efforts are being directed toward improving the efficiency of

Table 1. Background Literature: Surveys about GNNs (First Block) and Including GNNs (Second Block)

Study [Reference] (Year)	Contributions
Relational Inductive Biases, Deep Learning, and Graph Networks [11] (2018)	<ul style="list-style-type: none"> • Presents the idea of a graph network as a generalization of GNNs with building blocks • Encompasses well-known models, such as fully connected, convolutional and recurrent networks.
Graph Neural Networks: A Review of Methods and Applications [185] (2018)	<ul style="list-style-type: none"> • Presents a survey of the various GNN models • Discusses the applications where GNNs can be utilized and provides a taxonomy • Proposes open research problems, such as dynamicity and scalability in GNNs
A Comprehensive Survey on Graph Neural Networks [160] (2021)	<ul style="list-style-type: none"> • Overviews of GNNs in data mining and machine learning areas • Provisions a taxonomy for GNN models • Details the application areas of GNNs • Presents potential research directions, such as in scalability, dynamicity of GNNs, and so on.
Deep Learning on Graphs: A Survey [181] (2020)	<ul style="list-style-type: none"> • Provides a discussion on graph versions of recurrent and convolutional networks, autoencoders, reinforcement-learning and adversarial methods • Presents the application areas and future research directions for deep learning on graphs
Machine Learning on Graphs: A Model and Comprehensive Taxonomy [19] (2020)	<ul style="list-style-type: none"> • Presents a taxonomy to classify graph learning methods, from graph embeddings to GNNs • Proposes an encoder-decoder model that unifies all methods in a single approach • Expresses 30+ graph learning techniques using the proposed model
Graph Neural Networks Meet Neural-Symbolic Computing: A Survey and Perspective [91] (2020)	<ul style="list-style-type: none"> • Elaborates the relationship between GNNs and Neural-Symbolic Computing • Develops multiple GNN models with the perspective of being applied to Neural-Symbolic computing
Geometric Deep Learning: Going beyond Euclidean data [16] (2017)	<ul style="list-style-type: none"> • Proposes Geometric Deep Learning as an umbrella term for models that operate on non-euclidean dataset representations, including GNNs. • Within GNNs, provides a thorough review of convolutional models
Representation Learning on Graphs: Methods and Applications [66] (2017)	<ul style="list-style-type: none"> • Reviews the advancements in the area of representation learning on graphs • Primary focus is on the network embedding methods

algorithms, especially for large graphs, and toward demonstrating their efficacy for the aforementioned application areas. The interested reader will find multiple reviews of the state of the art in GNN algorithms and applications in the literature [11, 16, 19, 66, 91, 160, 181, 185], most of which we briefly analyze in Table 1. Other key aspects relevant or adjacent to GNNs such as network embedding [31], graph attention models [93], or network structure inference [17] have also received a comprehensive review.

As we will see along this article, however, less attention has been placed on the efficient processing of such new type of neural networks. While the issue has already been investigated in significant depth for CNNs or RNNs [24, 25, 39, 68, 90, 111], GNN processing remains largely unexplored.

Fig. 2. Graphical abstract of this survey from the GNN fundamentals (Section 2) to the proposed architectural vision (Section 5).

This is because GNNs are relatively novel and pose unique computing challenges, including the need to (i) support both dense and extremely sparse operations, (ii) adapt the computation to the specific GNN algorithm variant and the structure of the graph at hand, and (iii) scale to very large graphs to realize its potential in certain applications. Even though advances in sparse/irregular tensor processing [34] and graph processing [63, 154] may prove useful in accelerating GNNs, addressing their unique computing challenges requires more specialized proposals. Some attempts have been done from a software perspective, i.e., adapting the GNN operations to better match the capabilities of CPUs or GPUs [106, 144, 155]; and from a hardware perspective, i.e., designing custom processors tailored to the demands of GNNs [7, 53, 103, 164]. However, recent surveys and reviews [11, 16, 19, 66, 91, 160, 181, 185] lack a comprehensive analysis of such advances.

This article aims to bridge this gap by presenting, for the first time, a review of the field of GNNs from the perspective of computing. To that end, we make the following contributions as summarized in Figure 2: We start by providing a comprehensive and tutorial-like description of the fundamentals of GNNs, trying to unify notation. Then, using a **Knowledge Graph (KG)** approach, we chart the evolution of the field from its inception to the time of this writing, delving into the duality between GNN algorithms (seeing them as learning systems) and their associated computation (seeing them as sets of matrix multiplications and non-linear operations). From that analysis, we identify GNN computing as a nascent field. We finally focus on the computation aspect and provide an in-depth analysis of current software and hardware acceleration schemes, from which we also outline new potential research lines in GNN computing. To the best of the authors' knowledge, this is the first work providing a thorough review of GNN research from the perspective of computation, charting the evolution of the research area and analyzing existing libraries and accelerators.

The rest of this article is organized as follows: In Section 2, we discuss the basics of the GNNs. Section 3 presents the evolution of the research area from multiple perspectives. In Section 4, we expose the emergent area of GNN accelerators, summarizing recent works and elaborating upon the existing challenges and opportunities. Next, in Section 5, we present our vision for the architectural design of GNN accelerators with a focus on internal communication requirements. We conclude this article in Section 6.

2 FUNDAMENTALS OF GRAPH NEURAL NETWORKS

In this section, we discuss the basics of GNNs through a description of their building blocks and their role during the computation, both in inference and training.

Table 2. Graph Representation Notations

V	Set of vertices of the graph	$h_v, h_v^{(l)}, h_v^L$	Input, hidden, output feature vector of vertex v
E	Set of edges of the graph	$g_e, g_e^{(l)}, g_e^L$	Input, hidden, output feature vector of edge e
$N(v)$	Set of neighbours of vertex v	$\rho_V^{(l)}, \rho_E^{(l)}$	Node and edge aggregation functions of layer l
L	Number of GNN layers	$\phi_V^{(l)}, \phi_E^{(l)}$	Node and edge combination functions of layer l
y	Output global vector	$W_V^{(l)}, W_E^{(l)}$	Node and edge weight matrices of layer l

2.1 Notation

We first describe the main notation for GNNs as summarized in Table 2. Let a graph $G = (V, E)$ be defined by a set of vertices V , and a set of edges E that connect some of the vertices in V together. In particular, each vertex $v \in V$ has a neighbourhood set $N(v)$ determined by the edges connecting it to other vertices or the sampling set imposed by the GNN algorithm. Further, each vertex v contains a vertex feature representation h_v , and each edge $e \in E$ contains an edge feature representation g_e . The vertex or edge feature representations are generally one-dimensional vectors containing the scalar attributes that define them. Similarly, the graph may be associated to a global feature representation y containing graphwide attributes. For example, in a social networking graph, vertices might be users with attributes such as encoded name or location, whereas the edges might be the interaction between two users such as comments/likes on a picture. Graphwide features may be the number of users living a certain area or voting a certain political party.

GNNs essentially calculate a set of output feature representations for the vertices h_v , edges g_e , and complete graph y , respectively. Following with the example above, for targeting ads in a social network, output features of a vertex could be the probability of being interested in cars. It can thus be observed that, as in any other neural network, the dimensionality of the output feature vectors will be generally different than that of the input.

As we will see in Section 2.2, a GNN is divided in multiple layers. In each layer $l \in [1, L]$, there is an edge aggregation function $\rho_E^{(l)}$ and a node aggregation function $\rho_V^{(l)}$, as well as an edge combination function $\phi_E^{(l)}$ and a node combination function $\phi_V^{(l)}$. The combination functions may be neural networks involving matrices of weights $W_E^{(l)}$ and $W_V^{(l)}$ that are generally common to all edges and nodes, respectively. The outputs of an arbitrary intermediate layer l , given by its combination function, are hidden feature vectors $h_v^{(l)}$ and $g_e^{(l)}$. At the end of the GNN, besides obtaining the output node and edge feature vectors, h_v^L and g_e^L , there are global aggregation and combination functions ρ_G and ϕ_G , respectively, that provide final global output vector $\hat{y}$. Although most works assume that the graph is static, the computation may be repeated several times with evolving weight matrices to adapt to dynamic graphs [120].

We finally note that, due to the emergence of GNNs, aggregation and combination functions have taken different names in the literature. In an attempt to unify the notation, some equivalences are listed in Table 3.

2.2 General Structure

Fundamentally, a GNN is an algorithm that leverages the graph connectivity to learn and model the relationships between nodes. Through an iterative process that depends on the graph structure, the GNN takes the input edge, vertex, and graph feature vectors (representing their known attributes) and transforms them into output feature vectors (representing the target predictions). In general, the GNN operation contains the steps illustrated in Figure 3:

Table 3. Homogenized Nomenclature for Aggregate and Combine Functions in the Literature

Aggregation	Combination	Reference
Local transition	Local output	[130]
Aggregators		[65]
Aggregation	Update	[11]
Message + Aggregate	Update	[104]
Message	Update	[45, 56]
Message, reduce	Update	[151]
Scatter + ApplyEdge + Gather	ApplyVertex	[106]
Aggregation	Feature extraction + update	[103]
Gather + Reduce	Transform + Activate	[84]
Aggregation	DNN computation	[7]
Aggregation	Embedding	[53]
Aggregation	Combination	[163, 164]

- (1) **Pre-processing:** This is an initial and optional step generally done offline that can transform the input feature vectors and graph structure representation through a precoding process. This may be used to sample the graph, to re-order the graph toward reducing the algorithm complexity and its processing, or to encode the feature vectors, among others [23, 28, 65, 77, 141, 176, 181].
- (2) **Iterative updates:** After the pre-processing, the feature vectors of each edge and vertex are updated via the aggregate–combine functions iteratively. To update the edges, attributes from the edge itself, the connected vertices, and the graph are *aggregated* into a single set and *combined* to yield a new edge feature vector. Similarly, updating the vertices implies *aggregating* the feature vectors from neighboring vertices $N(v)$ and *combining* them to obtain a new feature vector. Note that each step or *layer* updates each edge and vertex with information coming from neighbours located at a single hop. Thus, the iterative process allows to gradually account for relations of increasingly distant nodes and edges. Additionally, in each successive layer, the graph may be coarsened by means of pooling [168] or the neighbourhood set changed by means of layer sampling [65].
- (3) **Decoding or readout:** If the graph has a global feature vector, then it is updated once after the edge and node updates are completed. The final output is either an edge/node embedding, which is a low-dimensional feature vector that represents edge- or node-specific information, or a graph embedding summarizing the information about the entire output graph instead.

As in any other neural network, the GNN processing depends on its architecture. GNNs are basically divided into *layers*, with each layer corresponding to one of the iterations in the update process described above. This means that each layer allows information from nodes to propagate further away from it. Hence, the precise number of required layers will depend on how relevant are the relations among *distant* nodes in a given application. The most widespread GNN algorithms have 1–5 layers [65, 87, 124, 146, 162] as an excessive amount of layers typically lead to the problems of feature oversmoothing, vanishing gradients, or overfitting [97]. A few works have proposed techniques to alleviate these issues and enable deep GNNs of up to 100 layers [22, 95], yet these proposals are in their infancy.

In each of the layers, information flows between vertices using an *aggregation* function and feature vectors are updated via the *combination* function after aggregation in a process similar to that of the classic **Weisfeiler-Lehman (WL)** test for graph isomorphism [157]. The size of

Fig. 3. GNN execution stages during inference: pre-coding, iterative process, and readout.

aggregation depends on the number of vertices and edges (ranging from hundreds to billions) whereas the size of combination depends on the length of the feature vectors (ranging from dozens of features to tens of thousands). The aggregation and combination functions for both edges and vertices are crucial design decisions as they determine the expressive power of the GNN, which has been demonstrated to be *at most* equal to the WL test in distinguishing graph structures [162]. As we will see in Section 3.2, Table 6, there is a wide variety of such functions ranging from simple averaging to weighted sums with learnable attention coefficients, different types of neural networks, from MLPs to LSTMs with their own weighted sums and non-linear activations, whose suitability depends on the relation to be learnt. The operations may vary across layers and differ between edges, vertices, or global updates. However, the structure is often simplified by (i) sharing the same operation across layers and (ii) removing or considering trivial combination functions for the updates of edges or nodes.

The fundamental structure here explained and depicted in Figure 3 can be complemented with sampling and pooling operations that help to reduce the computational complexity of GNNs [65, 168, 176], and/or augmented with support for dynamic graphs [120]. Sampling refers to the pruning of either the graph or the neighbourhood set of each node, and it is used to limit or harmonize the resources and runtime of the aggregation process, whereas pooling refers to the coarsening of the graph from one layer to the next, thus reducing the amount of nodes to process in both aggregation and combination. To add support for dynamic graphs, whose structure and input feature vectors may evolve over time, recurrent units are generally used to adapt the weight matrices in each timestep.

In summary, we can understand GNNs as a collection of neural networks working over a graph's connectivity. In the scope of each layer, we have up to two neural networks with learnable weights that determine the combination of edges and vertices, respectively. In the scope of the whole GNN, we have a neural network with learnable weights that determines the global update. The way these operations take place for inference and training is depicted next.

2.3 Computing GNN Inference

Algorithm 1 shows a pseudo-code describing GNN inference. The algorithm may take as inputs the feature vectors of the edges, vertices, and graph or initialize them. We can see how the execution is divided into layers (line 9) and, within each layer, each and every edge is updated in parallel by aggregating its own feature vector with those of the connected vertices (line 11). Each and every vertex is also updated in parallel by aggregating the feature vectors of its neighbours with

ALGORITHM 1: GNN Operations in Inference

```

1: procedure GNNOPERATOR
2:    $L \leftarrow$  Number of layers in the GNN
3:    $V \leftarrow$  Set of nodes in graph  $G$  (assumed static)
4:    $E \leftarrow$  Set of edges in graph  $G$  (assumed static)
   Initialize Nodes and Edges:
5:   for  $v \in V$  do
6:      $h_v^0 \leftarrow [x_v, 0, \dots, 0]$ 
7:   for  $e \in E$  do
8:      $g_e^0 \leftarrow [z_v, 0, \dots, 0]$ 
   GNN Layered processing:
9:   for  $l = 1$  to  $L$  do
     Edge processing: // Order of edge and node processing may be interchanged or even
       interspersed.
10:    for  $e \in E$  do // Order of aggregation and combination may be interchanged if
      aggregation is linear.
11:       $b_e^{(l)} = \rho_E^{(l)} (\{g_e^{(l-1)}, h_u^{(l-1)} : u \in N(e)\})$ 
12:       $g_e^{(l)} = \phi_E^{(l)} (\{b_e^{(l)}\})$ 
     Node processing: // Order of edge and node processing may be interchanged or even
       interspersed.
13:    for  $v \in V$  do // Order of aggregation and combination may be interchanged if
      aggregation is linear.
14:       $a_v^{(l)} = \rho_V^{(l)} (\{h_v^{(l-1)}, h_u^{(l-1)} : u \in N(v)\})$ 
15:       $h_v^{(l)} = \phi_V^{(l)} (\{a_v^{(l)}\})$ 
     Readout:
16:      $\hat{y} = \phi_G (\rho_G (\{h_v^L, g_e^L : v, e \in G\}))$ 

```

itself (line 15). The aggregated edges and vertices are transformed via combination functions (lines 13 and 17), which can be neural networks as we see in Section 3.2. Following the completion of the iterative process, a readout is performed using the corresponding function, which may again possibly be a neural network (line 18).

For an arbitrary layer $l \in [1, L]$, edge transformation occurs as

$$\text{AGGREGATION: } b_e^{(l)} = \rho_E^{(l)} (\{g_e^{(l-1)}, h_u^{(l-1)} : u \in N(e)\}), \quad (1)$$

$$\text{COMBINATION: } g_e^{(l)} = \phi_E^{(l)} (\{b_e^{(l)}\}), \quad (2)$$

so that the aggregation of edges ρ_E takes the feature vector g_e of the edge itself e , as well as the feature vectors of the vertices at its endpoints, h_u with $u \in N(e)$, for the previous layer $l - 1$. The combination ϕ_E uses this aggregation as input [162]. A similar reasoning applies to the aggregation and combination of vertices

$$\text{AGGREGATION: } a_v^{(l)} = \rho_V^{(l)} (\{h_v^{(l-1)}, h_u^{(l-1)} : u \in N(v)\}), \quad (3)$$

$$\text{COMBINATION: } h_v^{(l)} = \phi_V^{(l)} (\{a_v^{(l)}\}). \quad (4)$$

The equations describe how $a_v^{(l)}$ is calculated as the aggregation of the feature vectors from the nodes that are neighbours to v , from the previous layer $l - 1$, and how the feature vector of layer

Fig. 4. Example of computation in a sample GNN with node-level aggregation in inference (top left to top right) and training (bottom right to bottom left). The GNN has two layers, mean as the aggregation operator, and weighted ReLU for the combination. We show operations for node 1 only.

l is calculated using the aggregation $a_v^{(l)}$ as input. Last, a final readout function is applied, which may involve the aggregation and combination of feature vectors from edges and vertices of the entire graph, and from the last iteration L , hence obtaining the output feature vector $\hat{y}$ as

$$\text{READOUT: } \hat{y} = \phi_G \left(\rho_G \left(\{h_v^L, g_e^L : v, e \in G\} \right) \right). \quad (5)$$

Algorithm 1 hinges on the general assumption that aggregation and combination functions are (i) invariant to permutation of nodes and edges, since there does not exist any implicit order in a graph structure, unless some node feature indicates such an order, and (ii) invariant to the number of input nodes, since the degree of nodes may vary widely across the graph [11]. This implies that the functions within a layer can be applied to all edges and all vertices in parallel, following any order. Further, the order between aggregation and combination can be switched if the aggregation function is linear [103]. However, it is important that the order of layers is preserved to avoid violating data dependencies, which implies that all edge and node operations of layer l shall finish before starting those of $l + 1$.

To exemplify the computation occurring in inference, top charts of Figure 4 represent the layers of a simple GNN with vertex aggregation and combination only. We show the operations from the perspective of node 1, although all nodes would be realizing the same computations concurrently. We illustrate how the graph connectivity drives the aggregation from nodes 2, 3, and 6 into node 1, and that combination reduces the length of the feature vector through the weight matrices $W^{(1)}$. We note, however, that combination functions do not necessarily reduce the length of the feature vectors; that depends on the actual GNN architecture. The second layer repeats the exact same sequence, again reducing the length of the feature vector, this time through a different weight matrix $W^{(2)}$.

Extended notation for sampling, pooling, and dynamic graphs: As described above, sampling and pooling might impact the length aggregation and combination stages, whereas dealing with evolving graphs may require extra computation steps. Following the notation above, sampling essentially modifies either the input graph, G_s [176], or the neighbourhood operator making it dependent on the layer being computed $N_s^{(l)}(v)$. Pooling can be seen as a graph transformation across layers, thus making the set of nodes and edges to vary as well $E^{(l)}$, $V^{(l)}$. Finally, support

Table 4. Equivalence between General and MPNN Formulations

General	MPNN	Comments
l	t	Layer, timestep, or epoch
v	v	Node or vertex of interest
$u \in N(v)$	$w \in N(v)$	Node within the neighboring set $N(v)$ of node v
$h_u^{(l)}$	h_w^t	Feature vector of vertex u at layer l or epoch t
$\rho^{(l)}(\{h_u^{(l-1)} : u \in N(v)\})$	$\sum_{w \in N(v)} M_t(h_v^{t-1}, h_w^{t-1}, e_{vw})$	Aggregation at a layer or epoch with $M_t(\cdot)$ and $\rho^{(l)}(\cdot)$ as aggregation functions
$a_v^{(l)}$	m_v^t	Aggregated feature vector
$\phi^{(l)}(\{a_v^{(l)}\})$	$h_v^{t+1} = U_t(h_v^t, m_v^{t+1})$	Combination with functions $U_t(\cdot)$ and $\phi^{(l)}(\cdot)$ in a given layer or epoch

for dynamic graphs G_t requires the entire GNN to be time-dependent, introducing time in the notation. Neighbourhood sets, feature vectors, and most importantly, weight matrices would evolve over time, $N_t(\cdot)$, $h_{v,t}^{(l)}$, $g_{e,t}^{(l)}$, $W_{E,t}^{(l)}$, $W_{V,t}^{(l)}$.

Message passing equivalence: We note that notation relative to GNN algorithms is diverse in the literature. A notable example is that of **Message Passing Neural Network (MPNN)** [11], which describes the aggregations as message passing functions $M(\cdot)$, the combinations as update functions $U(\cdot)$, or the layers as time steps. Table 4 illustrates the equivalence between the MPNN formulations and the corresponding generic formulations from Equations (1)–(5).

Matrix multiplication notation: GNNs are typically expressed in matrix notation that helps understanding the underlying computation. An example for node classification with sum aggregation function is as follows. Let A be the normalized adjacency matrix of the input graph, $H^{(l)}$ the matrix of features for layer l , and $W^{(l)} = W_V^{(l)}$ the weight matrix for the vertex combination function. Then, the forward propagation to layer $l + 1$ can be expressed as

$$H^{(l+1)} = \sigma(AH^{(l)}W^{(l)}), \quad (6)$$

where $\sigma(\cdot)$ is the non-linear activation function, e.g., a ReLU. For more complex GNNs and aggregation-combination functions, the forward propagation equation may change.

2.4 Computing GNN Training

Aggregation, combination, and readout functions can be neural networks that may need to be trained before deployment. Training is performed via modifications of the traditional backpropagation algorithms, which take into account the unique traits of a GNN. Since a GNN unfolds into L layers similarly to a RNN, most GNNs employ **Back-Propagation-Through-Time (BPTT)** schemes or variants of it. A popular variant of BPTT is the Pineda-Almeida algorithm [5, 122], which relaxes the memory requirements as already mentioned in the seminal work by Scarselli et al. [130].

Specifically, in BPTT, a forward pass is first performed on the unfolded version of the GNN with its L layers. The loss function ε is then computed and the necessary gradient is backpropagated across layers. Since the weights are shared among all L layers, they are updated accordingly. This process is carried out recurrently with multiple samples, often grouped in batches, until some target accuracy is reached. Depending on the problem, a sample can refer to the entire graph (e.g., representing a specific molecule) or a portion of it (e.g., a set of users in a recommendation system).

To exemplify the computation occurring during training, bottom charts of Figure 4 represent backpropagation in a two-layer GNN. Again, we show the operations from the perspective of node

1, although all nodes would be realizing similar computations at the same time. The backward pass implies calculating the gradient of the loss function with respect to the weights first, via partial derivative over $W^{(2)}$, and then with respect to each vertex's feature vector. The operation is then repeated for the first layer, via its own weight matrix $W^{(1)}$ and each vertex's feature vector. The derivatives of the loss function are, eventually, used to update the weight matrices.

The computation of the loss function depends on the type of learning. While graph-centric approaches tend to be trained using supervised learning, node-centric approaches are usually trained by means of semi-supervised learning, wherein information of the node features from a portion of the graph, and not the whole graph, is utilized. An example of the former method can be learning if a specific new molecule (graph) has a certain property, using a GNN trained with molecules (graphs) whose properties are known beforehand and used as ground truth [56]. For the latter method, an example can be a recommender system. In such a system, a graph represents a store with nodes being shopping items and their features, and edges being relations among items. The output feature vector could describe how likely a given user will be satisfied with a particular item. In this case, a priori complete information is not available and semi-supervised learning from past purchases by this and other users (a subgraph) is used instead [167].

Matrix multiplication notation: To express backpropagation in a compact manner, we adapt the formulation in Reference [144] to the notation introduced in the previous section. Let $Z^{(l)} = AH^{(l)}W^{(l)}$ so that $H^{(l+1)} = \sigma(Z^{(l)})$. Then, the backpropagation starts by calculating the gradient of the loss function ε , which we denote as Y , with respect to the weight matrix of the last layer. For an arbitrary layer l , this operation yields

$$Y^{(l-1)} = \frac{\partial \varepsilon}{\partial W^{(l)}} = (H^{(l-1)})^T AG^{(l)}, \quad (7)$$

where $G^{(l)}$ is the gradient with respect to $Z^{(l)}$ and T denotes a transpose matrix. Therefore, $G^{(l)}$ refers to the propagation of the error back to each particular aggregated feature vector, yielding

$$G^{(l-1)} = AG^{(l)}(W^{(l)})^T \odot \sigma'(Z^{(l-1)}), \quad \text{with } G^L = \frac{\partial \varepsilon}{\partial Z^{(L)}} = \frac{\partial \varepsilon}{\partial H^{(L)}} \sigma'(Z^{(L)}), \quad (8)$$

where σ' is the derivative of the non-linear activation function.

3 THE EVOLUTION OF THE GNN FIELD

In this section, we aim to demonstrate that GNN computing is in an early yet rising stage as compared to the rest of GNN disciplines. We also observe that there is a wide variety of GNN algorithms that, as we will see in Section 4, complicate the task of designing accelerators. To these ends, we present the evolution of the body of knowledge in the area of GNNs from a general perspective in Section 3.1 and from an algorithm perspective in Section 3.2.

The study uses a KG approach that naturally exposes the confluence of multiple interrelated sub-fields in the GNN landscape. To generate the KG, a repository of annotated papers has been created. The papers are classified by their year of publication and are manually given a single tag using the title and keywords as main reference. Further, the references of each paper are extracted by means of the CERMINE library [142]. The generated database is introduced into the Neo4j graph tool [156], which allows to visualize the KG with nodes and edges representing papers and their citation relations, respectively. To highlight the category and importance of papers, vertices are color-coded depending on the paper category and sized proportionally to the number of citations.

Table 5. The Different Categories for the Classification of the State of the Art in GNNs

Tag Name	Meaning	Origins	References	Fraction
GNN modeling	This category includes the papers that encompass the topics of design and mathematical formulation of GNNs. Other salient design formalisms related to GNNs have been also categorized in this tag.	2005. While the most important paper in GNN modeling is from Scarselli et al. in 2009, it extends a seminal work from 2005. It defines the mathematical foundation of these GNNs, and thus becomes a fundamental paper in this category.	[40, 42, 57, 64, 65, 69, 79, 87, 118, 119, 130, 133]	13.19%
GNN applications	Papers with this tag elaborate upon the various applications of GNNs, regardless of the field.	2005. Given the ubiquity of graphs in real-world data, this is one of the first sub-fields to have emerged. In their seminal work, Scarselli et al. presented the first possible applications together with the first GNN model [132]. Since then, many other applications have appeared.	[12, 32, 41, 44, 47, 61, 83, 98, 108, 112, 114, 123, 126, 128, 129, 132, 136, 147, 167, 174, 180, 186]	24.17%
GNN complexity	This tag encompasses the papers that explore the complexity within the GNN structure and its operations.	2009. The exploration of the complexity of GNN execution may have started with Reference [131] in 2009, which analyzed the complexity for the most common GNNs at that moment. After this, we have to wait until 2017 to find more works that take into account complexity, as datasets become more resource demanding and large-scale applications become apparent.	[13, 18, 21, 26, 38, 82, 117, 122, 131, 148, 159, 169, 188]	14.29%
GNN algorithms	This tag refers to papers that introduce new algorithm variants to the GNN family, including aspects such as attention, isomorphism, sampling, or new operations at the aggregate–combine phases.	2009. We consider [130] as the first unification of multiple similar prior approaches. Others have attempted to do similar generalizations, such as the MPNN from Gilmer et al. [56] or the GN from Battaglia et al. [11].	[9, 11, 20, 35, 50, 56, 66, 86, 101, 102, 107, 113, 121, 125, 135, 140, 146, 152, 160, 162, 166, 179, 185]	25.27%
GNN accelerators	Under this tag, we gather papers that target the acceleration of GNNs either via software or hardware.	2017. The earliest paper to tackle the problem of GNN acceleration is Reference [65], in 2017, through a simplification of the algorithm via sampling. More recent works on software in CPUs and GPUs, and hardware acceleration in custom architectures, have also been considered.	[7, 23, 53, 85, 103, 155, 164, 175, 177, 187]	10.99%
GNN HW/SW requirements	This tag gathers works that, with the increasing popularity of GNNs as well as the complexity of the data-sets, analyzed the actual computational needs required to address these challenges.	2018. This specific sub-field started to gain traction in 2018, with the first work leading to Reference [106] where the hardware and software efficiencies in executing GNNs were studied.	[8, 76, 77, 106, 141, 144, 163, 182]	8.79%
GNN dataflow	Dataflow refers to the movement of data within the processing engine, which becomes crucial for the design of custom accelerators. Hence, under this tag we categorize the papers that formally describe possible dataflow solutions.	2018. Two primary works, i.e., Reference [106], which covers scalability in the training, and Reference [104], which covers efficiency for partitioning of the graph data, emerged.	[84, 104, 106]	3.30%

3.1 A General Perspective

Our first treatment of the GNN literature consists in classifying the papers by discipline. Concretely, we define the following taxonomy with topics ranging from formal mathematical aspects, to the algorithms, applications, and computing aspects: *GNN modeling*, *GNN applications*, *GNN complexity*, *GNN algorithms*, *GNN accelerators*, *GNN HW/SW requirements*, and *GNN dataflow*. The description of each topic, together with a discussion of its first works and the list of its references is given in Table 5. We also show the percentage of papers that pertain to a given category.

An important finding from our analysis is that the percentage of papers being categorized for *GNN accelerators*, *GNN HW/SW requirements*, and *GNN dataflow* are 10.99%, 8.79%, and 3.30%, respectively. These categories mostly relate to the computing side of GNNs as they concern the analysis of computational requirements of GNNs, optimization of GNNs via software, and development of hardware accelerators. We thus observe that a very small percentage of the existing research has approached GNNs from the perspective of computing. We further note that the first works to deal with these topics date back from 2017, when the very first specific paper on GNN acceleration was published. It can be therefore concluded that GNN processing is in its nascent

Fig. 5. Full knowledge graph representation as of October 2020.

stages of development. This is the main reason for computing aspects not being analyzed in depth in recent GNN surveys [11, 16, 19, 66, 91, 160, 181, 185], which we aim to address in this work, and also represents an opportunity to make an early impact in the GNN research field.

A second order analysis stems from the careful observation of the KG, which we show in Figure 5. In the left plot, the size of the node represents the aggregated number of papers in a category, whereas the thickness of the edge between two nodes illustrates the relative amount of citations between the papers of a given pair of categories. In the right plot, we can also analyze the connections between the papers within the same category. Several observations can be made:

- (i) The categories related to computing are small yet well connected to the theoretical side of GNNs, corroborating our earlier observation from Table 5.
- (ii) The algorithms sub-field is large as many papers have appeared implementing multiple variants in the heterogeneous group of methods that GNN is. We review the evolution of GNN algorithms later in Section 3.2.
- (iii) The applications sub-field is large but sparsely connected internally, which means that application papers are generally not aware of other applications, unless reusing some specific common mechanism. This may be due to the wide variety of application fields for GNNs, ranging from social networks to chemistry, computer networks, or even material science as analyzed in previous sections.
- (iv) The algorithm and application categories have a strong inter-connectivity, as each application paper shall at least mention the algorithms used to implement the proposed system.
- (v) The connection from application papers to computing papers is weak. This may be due to the relative immaturity of the GNN computing field and this may change in upcoming years, especially if applications clearly benefiting from specialized accelerators arise (akin to the appearance of CNN accelerators for computer vision).

To further understand the state of things in GNNs, we visualize the evolution of the field over time. Specifically, we plot the growth of the KG and of the amount of published papers over the

Fig. 6. Evolution of the GNN knowledge graph over the years 2009, 2012, 2015, 2018, and 2020 (with the color code from Figure 5) and cumulative number of papers published in GNN in general and computing in particular.

years in Figure 6. First works started to appear as soon as 2005 [57] and, at that point, most research efforts were centered around new algorithms and possible applications. Evolution was rather slow for a decade, which we attribute to the lack of a killer application and the modest popularity of deep learning methods at that time. The field exploded around 2016, when CNNs and RNNs were already well established. Such a dramatic growth coincides with the introduction of the **Graph Convolutional Networks (GCN)** [86], one of the first and most popular models for GNNs, later followed by the introduction of the message passing notation and quantum chemistry application in Reference [56]. We further observe that research on GNN computing started in 2017 and, since then, attained a similar growth to that of the field. This trend may be an indicator of a strong increase of related works in the near future. Hence, it can be concluded that the area of GNN accelerator design and development is emerging and, thus, necessitates deeper insights that we provide in upcoming sections.

3.2 An Algorithm Perspective

GNNs are a set of models with a vast amount of possible configurations and design decisions that allow to modulate the inductive bias of the algorithm. We have seen how, due to their flexibility and potential applicability, the family of GNN algorithms has grown rapidly in recent years. Since different algorithms may be more or less amenable to certain acceleration techniques, here we briefly summarize the progress in this sub-field from graph kernels to modern GNN algorithms. Note that a deep review of existing GNN algorithms is not the main focus of this work. For such an analysis, we refer the reader to more specific surveys [11, 19, 160, 181, 185].

Pre-GNN techniques. Prior to the advent of GNNs, relational information extraction from graphs was based on graph embeddings, i.e., the pre-processing of the graph to condense the information in a low-dimensional space thus making it amenable to traditional ML algorithms [16, 31]. Similarly, **Graph Kernels (GK)** are a family of methods that, after extracting graph-level embeddings

Fig. 7. GNN algorithm taxonomy based on model architectures and training strategies, adapted from [181] and [160].

of two or more graphs, compare them for classification tasks [55, 70]. An example of such approach is the random walk kernel, wherein random walks are performed on the graphs while simultaneously counting the matching walks [52]. As compared to GNNs, GKs are easier to train, because they have less hyperparameters, which on the other hand limits their performance. The main reason stems in the loss of potential information incurred by the process of graph embedding. Thus, to achieve acceptable performance, GKs require handcrafted (not learned) feature maps, whilst GNNs do not. GNNs retain the inherent graph structure as a powerful and expressive form of defining the neural network, instead of distilling the essence of the graph to feed a conventional neural network.

GNN algorithm classifications. Since the seminal work by Scarselli et al. [130], multiple approaches have been published with the aim of elaborating and complementing the GNN concept [6, 37, 69, 118] and many classifications can be carried out. A common distinction relates to the fundamental model upon which the GNN is built, for which a few taxonomies can be found in existing surveys [11, 19, 160, 181, 185]. As a reference, Figure 7 reproduces the classification made in Reference [185], which mostly differentiates between recurrent-based GNNs, convolutional-based GNNs, graph autoencoders, graph reinforcement learning, and graph adversarial networks. We added the remark made in Reference [160], where combinations of recurrent and convolutional approaches are termed as spatial-temporal.

On the one hand, recurrent-based GNNs refer to the initial GNN models including that of Scarselli [130], which employ recurrent units as the combination function. Other examples are CommNet [137], which operates over simple aggregations without edge transformations, or **Gated Graph Neural Networks (GG-NN)** [102], which use gated recurrent units [30] as the update function to improve convergence. On the other hand, convolutional-based GNNs expand the idea of convolution in the graph space [27] and can be divided into spectral-based [69] and spatial-based GNNs [186]. On the one hand, spectral-based models are built on spectral graph theory using graph signal processing techniques such as eigenvalue decomposition and filtering. However, they are computationally expensive methods, since the entire graph must be considered at once. On the other hand, spatial-based GNNs are much more computationally affordable, flexible, and scalable, since they only need to perform convolutions to the aggregation of features from neighbouring vertices [186]. Finally, spatial-temporal GNNs use both the spatial approach of the convolutions with the temporal approach of the recurrent units. An example is the network in **gated graph convolutional network (G-GCN)** from Reference [15].

Due to their flexibility and scalability, spatial-based convolutional GNNs are arguably the most popular model [20, 29, 48, 71, 99, 133, 143, 158, 165, 172]. In this paradigm, basic algorithms use a

mean function as aggregation, sometimes also taking the degree of neighboring into account [87], after which many variants followed. GraphSAGE incorporated information of self-node features from previous layers in the update function and also pioneered the concept of sampling in GNNs to reduce the computational cost of aggregation [65]. FastGCN [20] also uses the sampling idea and integrates other strategies to speed up computations, such as evaluating integral formulations using Monte Carlo sampling. Another simplifying operation is the differential pooling of DiffPool [168], which forms hierarchical clusters so that later layers operate on coarser graphs. On a different approach, **Graph Isomorphism Network (GIN)** [159, 162] proved that the conditions needed for a GNN to achieve the maximum expressive power in capturing the structure of a graph are to emulate a WL test [157]. The particularity occurs at the graph output feature vector, which is obtained by concatenating the readout vectors of all layers. We finally highlight **Graph Attention Networks (GAT)** as the enabler of multiple works in Natural Language Processing [145] and a particular case of the popular transformers approach. GATs update the node features through a pairwise function between the nodes with learnable weights [146]. This allows to operate with a learnt attention mechanism that describes the utility of the edges.

Another branch of GNNs are the so-called **Graph Autoencoders (GAE)** [86]. These GNNs are generative, which means that they convert a graph into a latent representation (i.e., encoding) that can be later expanded to generate to a new graph close in structure to the original one (i.e., decoding). What make these techniques unique in the graph domain is that GCNs may be used to generate the low-dimensional vectors in the encoding process [127]. GAEs are also typically trained using adversarial techniques, giving rise to graph adversarial networks such as NetRA [173].

We finally highlight that GNNs can be combined with reinforcement learning to give rise to novel graph learning techniques. For instance, MolGAN [35] generates molecular graphs with a certain end goal (reward). Another example is MINERVA, where reinforcement learning helps to predict the next node in the reasoning path of a KG [33].

Comprehensive frameworks. An aspect worth mentioning is that, within this multitude of algorithms, several groups have attempted to unify methods. One of the most popular ones is the message passing scheme [56, 183], whose operation and description are amenable to convolutional networks for learning molecular fingerprints [41], the classification methodology with GCN from [87], the interactive networks utilized for learning relationships and features [12], or also different flavours of Gated GNNs, to name a few. A further approach is that of the **Non-Local Neural Networks (NLNN)** [152] aimed at unifying various attention approaches including GATs. These generally do not include edges features or aggregations and, instead, just involve pairwise scalar attention weights between nodes. Both MPNN and NLNN are also included into a further approach to unification referred to as Graph Networks and proposed in Reference [11]. There, update functions applied to nodes, edges, or the complete graph are treated as differentiated blocks. The combination or repetition of several of these blocks gives rise to the different types of GNN found in the literature. Finally, Chami et al. propose an encoder-decoder model to express different graph embedding, graph regularization, graph auto-encoder, and GNN techniques [19].

Programming models. From the perspective of computation, several programming abstractions are considered to support all possible operations within any GNN, generally compatible with the aggregate-combine model. These models are useful when the matrix multiplication notation cannot be employed, because the aggregation or combination operations are not amenable to it or because the adjacency matrix is extremely sparse and suggests the use of other representations such as compressed sparse row or column. In fact, as we will see in the next section, multiple accelerators implement GNN-oriented programming models.

Among the different possible models, we highlight the **Scatter-ApplyEdge-Gather-ApplyVertex with Neural Networks (SAGA-NN)** from Reference [106], which is followed implicitly in most modern libraries [182]. SAGA-NN augments classical scatter-gather approaches with two operations and works as follows: Scatter sends the nodes' feature vectors through their edges and ApplyEdge performs edge combination with the scattered vectors. Then, Gather allows each vertex to aggregate the vectors from its neighbours, and ApplyVertex performs the vertex combination after the gather operation. Another proposed model is that of **Gather-Reduce-Transform-Activate (GReTA)** from Reference [84]. In this case, the four operations are user-defined and can be modified to implement any GNN. Aggregation is performed through gather and reduce, which allow each vertex to obtain the features from their neighbours and accumulate them into a single value. Combination is then performed through transform and activate, which typically do the matrix multiplication and non-linear activation of the aggregated data. More recently, Wang et al. proposed the *NeighborSelection-Aggregation-Update* model, which adds a flexible neighbor selection layer to the more conventional aggregate-update [150].

4 THE REVOLUTION OF GNN ACCELERATION

The optimization of ML algorithms and the building of custom hardware for high performance and efficiency has experienced an explosive growth in recent years [25, 67]. This has come shortly after academia and industry have unveiled the outstanding potential of DNN algorithms and their all-pervasive applicability. As evidenced in previous sections, the field of GNNs is arriving at a similar turning point. At the time of this writing, research in GNN methods is already extensive and keeps refining the algorithms and investigating new applications with high potential impact. Therefore, a key research aspect in the years ahead will be how to compute GNNs efficiently to realize their full potential.

GNN computing presents a set of unique challenges [163, 182] that have rendered existing libraries and hardware platforms inefficient, including:

- (i) **The existence of multiple GNN variants, which may include edge, vertex, and graph-wide updates**, with a variety of aggregation and combination functions as illustrated in Table 6, and possibly incorporating pooling and graph/layer sampling operations as well [28, 176]. These functions affect aspects such as the choice of operations to accelerate, the relative computational complexity of aggregation and combination, or the ordering constraints among them and across layers. Hence, instead of using a single general acceleration technique, GNN may require finding the right combination of techniques that works for a particular GNN variant.
- (ii) **The dependence of computation on the characteristics of the input graph** in terms of size, sparsity, clustering, or the length of the associated feature vectors. Graph connectivity may follow a power-law distribution, be evenly distributed, or be bipartite. Since the computation fundamentally depends on the input graph, decisions such as the use of dense or sparse logic, the dataflow to implement, the partitioning strategy, or the partitions' mapping and scheduling may need to be changed within and across graphs to maximize performance [51, 77, 141]. The challenge is, therefore, to develop accelerators that can dynamically adapt to the graph characteristics.
- (iii) **A unique combination of computing characteristics from deep learning and graph processing, leading to alternate execution patterns**. More specifically, combination often implies MLP-like operations over a dense weight matrix, which is generally computation-bound [138]. In contrast, aggregation involves, among other operations, fetching groups of vertices that often lead to irregular memory patterns [59]. Optimizations in

Table 6. Operations in Popular GNN Algorithms

Algorithm	Aggregation (a)	Combination (h^{l+1})
GCN [87]	$mean(N(h^l))$	$ReLU(W_l \cdot a)$
GIN [162]	$mean(N(h^l))$	$MLP(W \cdot ((1 + \epsilon^l) \cdot h^l + a))$
GS-mean [65]	$mean(N(h^l))$	$\sigma(W_l \cdot Concat(a, h^l))$
GS-max [65]	$max_{j \in N(h^l)} \sigma(W_l^1 \cdot h_j^l)$	$\sigma(W_l^2 \cdot Concat(a, h^l))$
GS-LSTM [65]	$LSTM(rand(N(h^l)))$	$\sigma(W_l \cdot Concat(a, h^l))$
GAT [146]	$\sum_{j \in N(h^l)} \alpha_j W h_j^l$	$\sigma(a)$
HighwayGCN [124]	$\sigma(W^l \cdot h^l + b^l)$	$h^{l+1} \odot a + h^l \odot (1 - a)$
GRN [103]	$mean(N(h^l))$	$GRU(h^l, W^l \cdot a)$

Notation: σ is a nonlinear function, α_j is the attention coefficient, b is the bias, $\odot$ is a dot-product, *Concat* is matrix concatenation, *MLP* is a multi-layer perceptron, *GRU* a gated recurrent unit, and *LSTM* is Long short-term memory.

aggregation can be done via sparse GEMM of the adjacency matrix [163], but they are not generalizable to all graphs/GNNs and typically not enough to combat the extreme sparsity of adjacency matrices. Therefore, the challenge is to develop architectures that accelerate such distinct phases and their intertwining at runtime.

- (iv) **A wide pool of applications with not only different graph characteristics, but also different performance targets.** For example, recommendation systems need to scale to extremely large graphs of up to billions of edges and target high computational throughput. In contrast, applications such as object detection in point clouds [134] or fraud detection [153] rather need to focus on latency and energy efficiency. This highlights the need for acceleration techniques that address not only the challenging GNN computation at relatively small scales and in real time, but also the storage and multi-GPU coordination issues at larger scales.

A direct consequence of the aforementioned aspects is that the bottleneck or the critical operation/kernel may vary across GNNs or applications, as shown in References [10, 163, 182]. In light of these challenges, GNNs call for new solutions both in software and hardware. On the software side, several libraries have been proposed to improve the support for GNNs and efficiently compute its multiple variants both in inference and training. The extensions of popular libraries such as PyTorch or **Tensorflow (TF)** [1, 45, 58] are clear examples of this. On the hardware side, new accelerator architectures have been surfacing recently [53, 85, 103] that attempt to deal with the flexibility and scalability challenges of GNNs mostly in inference thus far. In the next subsections, we provide an exhaustive overview of existing techniques.

4.1 Software Frameworks and Accelerators

The challenges of GNN processing rendered both traditional DNN libraries and graph processing frameworks [63, 154] inefficient. The reason is the alternating computing phases of GNNs. DNN libraries would be good at speeding up combination operations within vertices and edges, but perform poorly during aggregation. Graph processing libraries, instead, do a good job at managing irregular memory accesses when traversing the graph. However, these libraries assume trivial operations at the vertices, which is not the case in GNNs. To bridge this gap, recent works have started investigating how to adapt the libraries to (i) provide easy to program interfaces to implement multiple GNN variants, (ii) handle the variety of potentially sparse GNN operations efficiently in widespread GPU hardware, and (iii) scale computations to large-scale graphs and multiple GPUs.

In the following, we review a comprehensive selection of software frameworks and accelerators, listed in Table 7. The analysis does not include GunRock [154] or GE-SpMM [74] for different reasons. GunRock, despite implementing GraphSAGE in its latest versions, is a graph processing library that does not exploit intra-vertex parallelism. In fact, two works detailed below [73, 155] achieve speedups of 30–200× with respect to GunRock. GE-SpMM, although claiming to be tailored to GNNs, is an acceleration method for general-purpose sparse matrix multiplication in GPUs.

A first observation from Table 7 is that software frameworks have been tested for a wide variety of GNN algorithms and relevant datasets. Around 20 different GNN variants have been evaluated, being GCN, GS, and GIN the most common. Even though Amazon, Reddit, Protein, Cora, or CiteSeer datasets are popular in the community, a lack of a widely adopted benchmark suite [72] makes the datasets to vary widely. It is worth noting, however, that graphs can range from hundreds of edges in chemistry applications to billions of edges in large-scale recommendation systems. As we see next, performance comparisons are scarce, but generally take **PyTorch Geometric (PyG)**, **TF**, and **Deep Graph Library (DGL)** as baselines and often report between one and two orders of magnitude improvement typically in CPU+GPU platforms, with some exceptions on multi-GPU systems [76, 106] or distributed computing clusters with up to 32K cores [150, 178]. Most of the tested frameworks provide optimizations that could work for both acceleration of both training and inference, yet the evaluation is unequal. Training is evaluated in References [45, 73, 76, 105, 106, 141, 150, 151, 178, 187], whereas inference time is only measured in References [73, 76, 77, 84, 155, 178].

PyTorch Geometric. PyG [45] is a widespread library that is built upon PyTorch and that provides support for relational learning, illustrated in a myriad of algorithms. The key aspect is the definition of a message passing interface with definition of message and update functions for neighbourhood aggregation and combination, respectively, and multiple pooling operations. To accelerate GNN processing, PyG handles sparsity via dedicated GPU scatter and gather kernels that operate in all edges and nodes in parallel, instead of using sparse matrix multiplication kernels. Relevantly, Facebook released Pytorch-BigGraph [94], a library that allows to process arbitrarily large graphs by introducing partitioning and distributed processing and that could complement PyG.

Deep Graph Library. DGL [151] is a recent library that works on top of TF, PyTorch, or MXNet, and provides plenty of examples and code for multiple GNNs. The library defines three functions: message for edge aggregation and update and reduce and update for aggregation and combination at the nodes. To boost performance, DGL takes a matrix multiplication approach and leverages specialized kernels for GPUs or TPUs. In particular, both sampled dense-dense and sparse matrix multiplications are considered together with node, edge or feature parallelization. As discussed in their work [151], DGL uses heuristics to choose among the different options as the optimal parallelization scheme depends on multiple factors including the input graph. Thanks to this approach, DGL claims to achieve an order of magnitude faster training than PyG. Recently, researchers at Amazon have released a DistDGL, a system based on DGL for distributed mini-batch training scalable to billion-edge graphs [184]. To achieve it, DistDGL uses min-cut graph partitioning via a lightweight algorithm.

NeuGraph. Microsoft Research led one of the first specialized frameworks for parallel processing of GNNs in GPUs, NeuGraph [106]. Although it is built on top of TF, NeuGraph is not open source at the time of this writing. The framework implements a programming model, SAGA-NN, based on the functions Scatter for edge aggregation, ApplyEdge for edge combination, Gather for node aggregation, and ApplyVertex for node combination. Scatter-gather kernels are used in the

Table 7. State of the Art in Software Frameworks and Accelerators for GNNs (GS = GraphSAGE)

Name	Main Features	Evaluation		
		Algorithms	Datasets	Baselines
PyG [45]	<ul style="list-style-type: none"> • Leverages widespread adoption of PyTorch. • Wide variety of example codes available. • Use of scatter-gather kernels + node/edge parallelism. • Evaluated in GPU. Compatible with BigGraph [94] to scale. 	GCN, GAT, SGC, GS, GIN, and so on...	Cora, CiteSeer, PubMed, MUTAG, Proteins, Collab, IMDB, Reddit	DGL
DGL [151]	<ul style="list-style-type: none"> • Library compatible with TF, PyTorch and MXNet. • Deep documentation and support, tutorials. • Based on matrix-mul kernels. Evaluation in CPU and GPU. • Augmented with DistDGL [184] for distributed computing. 	GCN, GAT, SGC, GS, GIN, R-GCN, GCMC	Reddit, OGB (Arxiv, Protein, Product, Citation, PPA), MovieLens	PyG
NeuGraph [106]	<ul style="list-style-type: none"> • Implementation and evaluation for scaling to multiple GPUs. • Four-function model allowing for updates at edges and nodes. • Optimized partitioning, scheduling, pipelining, transfers. • Built on TF, not open sourced. 	GCN, CommNet, GG-NN	Pubmed, Blog, Reddit, Enwiki, Amazon	DGL, TF
AliGraph [187]	<ul style="list-style-type: none"> • Targeting large-scale graphs and distributed systems. • Emphasis on distributed storage and partitioning • Only work with heterogeneous and dynamic GNNs, and huge datasets (up to 483M edges, 6.5B edges). Built on top of TF. 	GS, six in-house algorithms	Amazon, Taobao	N/A
FlexGraph [150]	<ul style="list-style-type: none"> • Uses NAU programming model for flexible aggregation. • Hierarchical aggregation with dynamic sparse-dense logic. • Supports distributed computing, tested in 1500-core system. 	GCN, PinSage, MAGNN	Reddit, FB91, Twitter, IMDB	PyG, DGL, DistDGL, Euler
AGL [178]	<ul style="list-style-type: none"> • Aiming for scalability, fault tolerance, and integrality. • Uses MapReduce to scale, tested in 32000-core system. 	GCN, GS, GAT	Cora, PPI, UUG,	PyG, DGL
ROC [76]	<ul style="list-style-type: none"> • Implemented on top of FlexFlow [78]. • Optimizations: dynamic partitioning, memory management. • Evaluated with single and multiple GPUs via NVLink. 	GCN, GS, CommNet, GIN, FastGCN	Pubmed, PPI, Reddit, Amazon	TF, DGL, PyG, NeuGraph
GNN Advisor [155]	<ul style="list-style-type: none"> • Unique runtime profiling of graph information (degree, feature size, communities) to guide GPU processing • Extensive comparison with similar frameworks in single GPU 	GCN, GIN	CiteSeer, Cora, Pubmed, PPI, Prot, Yeast, DD,twit, SW620H, amazon, artist	DGL, PyG, GunRock, NeuGraph
PCGCN [141]	<ul style="list-style-type: none"> • Motivated by power-law distribution of node degrees. • Optimized partitioning to generate dense matrices. • Dual execution mode depending on sparsity of each partition. • Built on top of TF, evaluated in single GPU. 	GCN	Pubmed, Blog, Youtube, C1000-9, MANN-a81, Reddit, synthetic (RMAT)	TF, DGL, PyG
HAG [77]	<ul style="list-style-type: none"> • Removes redundant sums in aggregation by <i>fusing</i> nodes. • Runtime algorithm to <i>fuse</i> nodes only if predicted beneficial. • The impact on operation reduction is independent of hardware, but the impact on execution speed is not. 	GCN, GIN, SGC	BZR, PPI, Reddit, IMDB, COLLAB	N/A
FeatGraph [73]	<ul style="list-style-type: none"> • Optimized matmul kernels for aggregation and combination. • User-defined combination functions and optimizations. 	GCN, GS, GAT	OGB (Proteins), Reedit, sythetic graphs	GunRock
G ³ [105]	<ul style="list-style-type: none"> • Brings together graph processing frameworks and GNNs. • Offers APIs over C/C++ for ease of programming. • Uses GunRock [154] to provide GPU runtime optimizations. 	GCN, SGC	PubMed, Reddit	PyG, TF
GR _E TA [84]	<ul style="list-style-type: none"> • Programming abstraction with user-defined functions, similar to SAGA, targeting accelerators and any GNN variant. • Evaluation based on GRIP (see Table 8) in ASIC. 	GCN, GS, G-GCN, GIN	Youtube, Livejournal, Pokec, Reddit	N/A

functions of the same name, whereas matrix multiplication primitives are used in the combination functions. NeuGraph also features a number of optimizations to accelerate GNN computing. First, the partitioning of large graphs performed via the Kernighan-Lin algorithm to make partitions denser and minimize the transfers between partitions, which harm performance. Second, scheduling of partitions to the GPU is optimized by batching together small sparse partitions that can be computed together [115], and also profiling transfer and computation times in first GNN layer to later pipeline different chunks perfectly. Third, NeuGraph also eliminates redundant computation by fusing multiple edges together. Finally, it allow to scale GNN to multiple GPUs by distributing the computation, and optimizes the transfer of information by using a ring-based dataflow that minimizes contention at the interconnect.

AliGraph. Developed by the Alibaba group and open-sourced with the name of graph-learn, AliGraph is a GNN framework built on top of TF [187]. The framework is thought for the processing of very large and dynamic graphs in large-scale computing systems, and is currently used

in recommendation services at Alibaba. It implements three layers, namely *storage*, which implements partitioning with four different algorithms, but in this case to store the graph in a distributed way; *sampling*, which, unlike other frameworks, allows us to define custom sampling of a nodes' neighbourhood relevant to algorithms such as GraphSAGE; and *operator*, which implements the aggregation and combination functions. In overall, the AliGraph is unique due to its distributed approach and the many optimizations made at the storage layer to minimize data movement, such as the use of four different partitioning algorithms depending on the characteristics of the graph, or caching important vertices in multiple machines to reduce long misses.

FlexGraph. The Alibaba group also leads the development of FlexGraph [150], a distributed framework for GNN training whose distinct features are their flexible definitions of neighbourhood and the hierarchical aggregation schemes. To this end, FlexGraph uses the NAU programming model described in Section 3.2. To speedup training, FlexGraph combines hierarchical aggregation with a hybrid execution strategy combining sparse and dense logic. It also accelerates distributed execution through an application-driven workload balancing strategy and a pipeline processing strategy to overlap computations and communications.

AGL. AGL [178] is a framework created specifically for industrial deployments of massive GNNs. To that end, the authors emphasize their scalability, fault tolerance, and use of existing widespread methods for distributing the computation. In particular, AGL uses MapReduce [36] to that end and tests the proposed system in CPU clusters. The framework has three modules: one for creating independent neighbourhoods that can be processed in parallel, one for optimizing training, and one for the slicing of the graph and calculation of inference. Numerous optimizations are proposed in the sampling and indexing of the graph, partitioning and pruning, and pipelining of computation during training.

ROC. ROC [76] is another GNN framework targeting multi-GPU systems, in this case built on top of FlexFlow [78]. Similarly to AliGraph or AGL, ROC is able to distribute large graphs to multiple machines. However, this framework differs from others in that the partitioning method and memory management is performed with dynamic methods providing extra acceleration. First, ROC uses an online linear regression model to approach partitioning optimally. This model uses the training iterations to learn the best strategy of a specific graph, outperforming static methods significantly. Second, memory management is treated as a cost minimization problem and solved via an online algorithm that finds where to best store each partition. The authors demonstrate that such acceleration methods provide better scalability than DGL and PyG in single GPUs, and better scaling to multiple GPUs than NeuGraph.

GNNAdvisor. The work by Wang et al. [155] presents a runtime system that aims to systematically accelerate GNNs on GPUs. Instead of treating this problem via abstract models as done in ROC, GNNAdvisor does an online profiling of the input graph and GNN operations to guide the memory and workload management agents at the GPU. In particular, it leverages (i) the node degree to fine-tune the group-based workload management of the GPU, (ii) the size of the node embedding to optimize workload sharing, and (iii) the existing of communities within the graph to guide partitioning and scheduling. While the two first features are trivial to obtain, community detection is generally harder. In this case, the authors use a combination of node renumbering and Reverse Cuthill–McKee algorithm to reorder the adjacency matrix in a way that dense partitions are available. Thanks to all these techniques, the authors claim 3–4× speedup over DGL, PyG, and NeuGraph in a high-end GPU.

PCGCN. The paper by Tian and co-authors [141] present a partition-centric approach to acceleration of GNNs in GPUs, which they implement on top of TF. The contribution is motivated by the power-law distribution of the node degrees in a graph, which largely affects partitioning. PCGCN applies a locality-aware partitioning, METIS [81], that helps obtaining dense sub-matrices. That, however, does not prevent sparse partitions to appear. To combat this, PCGCN profiles the partitions at runtime and applies a dual-mode of operation: dense matrix representation and multiplication kernels when dense, and column-sparse representation and sparse kernels otherwise. In the paper, the authors compare their implementation with vanilla TF, and also DGL and PyG, and report the lowest speedup across libraries. Even in this case, PCGCN always speeds up execution and achieves up to 8.8× in highly clustered graphs.

HAG. This work presents the concept of **Hierarchically Aggregated computation Graph (HAG)** [77]. The authors make the observation that many of the operations made during the aggregation stage are repeated multiple times when nodes share similar neighbourhoods. In response to this, HAGs are presented as an alternative representation that proactively “fuses” nodes with common neighbourhoods, removing redundant aggregations during the execution of any GNN. Since the search of similarly connected nodes can be expensive, HAG employs a cost function to estimate the cost of certain node fusions, to then adopt a search algorithm affordable for runtime. With only 0.1% of memory overhead, HAG reduces the amount of aggregations by 6.3×.

FeatGraph. Developed in collaboration with Amazon, FeatGraph [73] proposes to optimize kernels of aggregation and combination separately. Different from other frameworks, here the user can define the combination function and ways to parallelize it, so that the scheduler can take it into account. As optimizations, FeatGraph also proposes to combine graph partitioning with feature dimension tiling and to adopt a hybrid partitioning scheme for GPUs.

G³. Liu et al. [105] propose a framework for the training of GNNs in GPU systems. G³ facilitates the task of GNN creation by providing a set of flexible APIs over C/C++ code that implement widespread layers and models. G³ also incorporates a set of graph-centric optimizations based on GunRock for aggregation [154] dealing with memory management, workload mapping, and load balancing. In training, G³ shows up to 100× speedup over PyG and TF in a high-end GPU.

GReTA GReTA [84] is a processing abstraction for GNNs aiming at simplifying their representation for hardware implementations. To this end, GReTA consists of four user-defined functions: Gather and Reduce to describe the aggregation, and Transform and Activate to describe the combination. These functions enable certain flexibility to accommodate different GNN types. GReTA also discusses partitioning briefly and exemplifies it in a hardware accelerator called GRIP [85], which is described in the next section.

Paddle Graph Learning. Developed by Baidu Research, Paddle Graph Learning [3] is a graph learning framework based on PaddlePaddle [109] that supports both walk-based and message passing models in heterogeneous graphs. Moreover, it integrates a Model Zoo supporting many GNN models to foster adoption, as well as support for distributed computing.

Tripathy et al. In this work, the authors compare multiple parallelization algorithms that partition and distribute the GNN in multiple GPU clusters, i.e., 1D, 1.5D, 2D, and 3D algorithms, and model the tradeoff between inter-GPU communication and memory requirements of these setups analytically and for training. The model takes a large adjacency matrix and breaks it down to a fixed amount of processes depending on the algorithm. Then, an analysis is made on the amount of effectual operations and results to be communicated across the GPUs. Their implementation over

Fig. 8. Qualitative classification and schematic representation of hardware accelerators for GNN inference. Green, blue, and red squares represent processors, memory, and control units, respectively.

PyG shows promising scalability and nominates the 1.5D algorithm as a promising and balanced alternative, although the best algorithm depends on the characteristics of the input graph.

4.2 Hardware Accelerators

We have seen above that software accelerators streamline the execution of GNNs in CPU-GPU platforms present in most computing systems, achieving significant speedups both in inference and training. Fewer works [8, 182] have tested GNN training in the TPUs typically used in *dense* DNNs, showing similar performance than in GPUs.

In this context, a pertinent question is whether custom hardware accelerators can tackle the unique challenges of GNN computing and live up to the promise of order-of-magnitude improvements that, to cite an example, have been already achieved in CNNs [25]. Pursuing this goal, several hardware accelerators have emerged that attempt to handle the extreme density and alternating computing requirements of GNNs. We next discuss all the designs published to date, using as reference the schematic diagrams of their architecture shown in Figure 8. The figure also tries to classify the architectures in two axes: unified versus tiled (to assess whether the computing phases are physically separated and how tightly coupled they are) and general (to specific to assess how easy is to adapt the accelerator to multiple GNN variants).

A summary of the main features of the accelerators and evaluated algorithms and datasets is given in Table 8. We observe that most works revolve around the GCN algorithm, which is popular and easy to illustrate. Datasets are generally smaller than in software acceleration works, mainly because of the memory limitations of hardware accelerators in inference and the cost of simulating hardware architectures. Cora, CiteSeer, and Reddit are the most common ones. While performance comparisons are difficult due to the many variables involved, most works use CPUs and GPUs as baselines and, in some cases, even HyGCN [164] and **Autotuning-Workload-Balancing GCN**

Table 8. State of the Art in Hardware Accelerators for GNNs

Name	Main Features	Evaluation		
		Algorithms	Datasets	Baselines
EnGN [103]	<ul style="list-style-type: none"> Unified architecture with dense hardware, single dataflow, generalizable to many GNN variants. Aggregation via RER. Optimizations: edge reordering, degree-aware vertex cache, scheduling. 	GCN, GS, GG-NN, GRN, R-GCN	Cora, PubMed, Nell, Reddit, Enwiki, Amazon, synthetic (RMAT), AIFB, MUTAG, BGS, AM	CPU-DGL, GPU-DGL, CPU-PyG, GPU-PyG, HyGCN
HyGCN [164]	<ul style="list-style-type: none"> Hybrid architecture with separate aggregate/combine phases. Fine-grained pipelining via inter-phase coordinator. Eliminates sparsity with window sliding/shrinking approach. Focused on GCNs, unclear how to generalize (no edge updates). 	GCN, GSC, GIN, DiffPool	IMDB, Cora, CiteSeer, COLLAB, PubMed, Reddit	CPU-PyG, GPU-PyG
AWB-GCN [53]	<ul style="list-style-type: none"> Adapts to varying GNN workloads via three load balancing techniques, chosen based on the sparsity of each partition. Processes combination first to reduce the number of operations. Fine-grained pipelining of aggregation and combination. Focused on GCNs, unclear how to generalize. 	GCN	Cora, CiteSeer, PubMed, Reddit, Nell	CPU-PyG, GPU-PyG, FPGA, HyGCN
GRIP [85]	<ul style="list-style-type: none"> Uses the GReTA abstraction [84], generalizable to any GNN. Actual implementation with techniques similar to HyGCN. 	GCN, GIN, G-GCN, GS	Youtube, Livejournal, Pokec, Reddit	CPU-TF, GPU-TF, TPU, HyGCN
Auten et al. [7]	<ul style="list-style-type: none"> Tiled architecture, ready for scale-out via Network-on-Chip. Similar to HyGCN, less specialized but easier to generalize. 	GCN, MPNN, GAT, PGNN	Cora, CiteSeer, DBLP, PubMed, QM9_1000	CPU, GPU
Zhang et al. [177]	<ul style="list-style-type: none"> Combination of offline software acceleration (redundancy elimination + node reordering) and hardware acceleration in FPGA. Optimizations: double buffering, node+feature parallelism, dual pipelining mode depending of order of matrix multiplications. 	GCN	Flickr, Reddit, Yelp	CPU-TF, GPU-TF, CPU-C++, GPU-C++
Rubik [23]	<ul style="list-style-type: none"> Hierarchical and unified PE array design Includes small caches to eliminate redundant aggregations Adds graph reordering in software to improve cache utilization 	GIN, GS	Collab, BZR, IMDB, DD, CiteSeer, Reddit	Eyeriss-like, GPU-PyG
GCNAX [96]	<ul style="list-style-type: none"> Architecture with reconfigurable loop ordering and fusion. Choice is made after an offline design space exploration. Uses outer product to mitigate unbalanced presence of zeros. 	GCN	Cora, CiteSeer, Pubmed, Nell, Reddit	HyGCN, AWB-GCN, SpArch
GraphACT [175]	<ul style="list-style-type: none"> Only accelerator evaluating training and memory footprint. CPU+FPGA. Optimizations rely on load balancing, scheduling, batching, removal of redundant aggregation operations. 	GCN	PPI, Reddit, Yelp	CPU, GPU

(AWB-GCN) [53] as early works on hardware acceleration. In general, the proposed accelerators are around two and three orders of magnitude faster and more energy efficient than GPU and CPU platforms, respectively, often occupying less than 10 mm². There is no consensus on which software framework shall be used in the baselines. Finally, all accelerator proposals except GraphACT are designed and evaluated for inference.

EnGN. Among the first accelerators to appear, EnGN [103] presents a unified architecture heavily inspired by CNN accelerators. The GNN is fundamentally treated as concatenated matrix multiplication of feature vectors, adjacency matrices, and weights—all scheduled in a single dataflow. An array of clustered **Processing Elements (PEs)** is fed by independent banks for the features, edges, and weights to compute the combination function. To perform the aggregation, each column of PEs is interconnected through a ring and results are passed along and added according to the adjacency matrix in a process the authors call **Ring-Edge Reduce (RER)**. Within this architecture, sparsity is handled with several optimizations. First, the RER aggregation may lead to multiple ineffectual computations for sparsely connected nodes. To avoid this, EnGN reorders edges on the fly in each step of the RER. Second, PE clusters are attached to a degree-aware vertex cache that holds data regarding high-degree vertices. The reasoning is that well-connected vertices will appear multiple times during the computation and caching them will provide high benefit at modest cost. Other optimized design decisions relate to the order of the matrix multiplications when the

aggregation function is sum, which affects the total number of operations, or the tiling strategy, which affects data reuse and I/O cost.

HyGCN. The authors of HyGCN [164] build upon the observation that GNNs present two main alternating phases of opposed computation needs and introduce a hybrid architecture for GCNs. HyGCN is composed of separate dedicated engines for the aggregation and the combination stages, plus a control mechanism that coordinates the pipelined execution of both functions. Being dense, the combination stage is computed via a conventional systolic array approach. The aggregation stage has a more elaborated architecture featuring a sampler, an edge scheduler, and a sparsity eliminator that feeds a set of SIMD cores. Within this architecture, sparsity is handled at the aggregation engine thanks to efficient scheduling and the sparsity eliminator. The latter takes a window-based sliding and shrinking approach to dynamically adapt to varying degrees of sparse multiplications. To further adapt to the workloads, HyGCN allows to group the SIMD cores in aggregation and the PEs in combination in different ways depending on the size of feature vectors. Finally, special attention is placed to the design of the inter-engine coordinator to optimize memory accesses and allow fine-grained pipelining of the execution toward maximizing parallelism dynamically.

AWB-GCN. The AWB-GCN accelerator [53] advocates for an aggressive adaptation to the structural sparsity of the GNN. The authors motivate their design by analyzing the power-law distribution of most graphs, arguing that some parts of the computation will be dense and others extraordinarily sparse, creating unbalances. To address the imbalance, the architecture develops a custom matrix multiplication engine with efficient support of skipping zeros. To that end, data from memory is fed via a **task distributor and queue (TDQ)** to a set of PEs and accumulators. The TDQ takes two designs adapted to when sparsity is moderate or high. Since AWB-GCN focuses on GCNs that have linear aggregation functions, the authors propose to process combination first as this generally reduces the amount of features and, thus, the amount of operations performed in aggregation. Furthermore, AWB-GCN provides a fine-grained pipelining mechanism to overlap the execution of combination and aggregation even within the same layer. However, the key of AWB-GCN are its three workload balancing functions. The first one is local and tries to balance the load among neighboring PEs. The second one is remote and attempts to pour overflowing computation from a busy PE to a single remote underutilized PE. The third one takes the load of extremely busy PEs processing very dense node clusters and divides across multiple idle PEs. To support that, AWB-GCN provisions hardware at the TDQ and the connections to the PEs to allow the remapping of nodes to remote PEs and to take them back for coherent aggregation. Moreover, all decisions are taken based on information extracted from simple counting at the queues.

GRIP. A key aspect of most existing accelerators is that they focus on GCNs as a relevant GNN algorithm. In contrast, the GRIP accelerator [85] leverages the abstraction of GReTA [84] to develop a general accelerator for any GNN variant, allowing to perform edge and node updates with user-defined functions. The GRIP architecture reflects this by having separated and custom units and accumulators for both edges (gather, reduce) and vertices (transform, activate). A control unit orchestrates data movement between the different units and respective buffers. In the sample implementation, GRIP divides the edge update unit into lanes to execute vertices simultaneously and takes an input-stationary dataflow for the vertex update unit. Among the optimizations made, we found pipelining and tiling adapted to the particularities of the implemented dataflows, similar to that of other accelerators.

Auten et al. Unlike most other accelerators, this work [7] proposes a modular architecture for convolutional GNNs. The basic unit of the accelerator is a tile composed by an **aggregator module (AGG)**, a **DNN accelerator module (DNA)**, a **DNN queue (DNQ)**, and a **graph PE (GPE)**,

all of them connected to an on-chip router. Thus, the architecture can be scaled out by interconnecting multiple tiles among them and with memory. Within each tile, the architecture has a similar structure than HyGCN, with the DNA being an array for dense multiplication, the AGG an edge-controlled adder, the DNQ taking the role of inter-engine buffer, and the GPE controlling execution. In this case, however, the GPE is a lightweight CPU managing multiple threads rather than an optimized controller.

Zhang et al. The work by Zhang and co-authors [177] presents a combination of software and hardware acceleration for GCNs. On the one hand, the graph is pre-processed via a redundancy elimination mechanism similar to that of Reference [77] and a node reordering similar to that of Reference [155]. Pre-processing is done offline and is justified for the repeated benefits that it can provide to multiple inferences to static graphs. The processed graph is then fed to a hardware accelerator implemented in a FPGA consisting of differentiated pipelined modules for aggregation (sparse array) and combination (dense systolic array and non-linear activation module). As differentiating elements with respect to other designs, we find that the aggregator module uses a double-buffering technique to hide latency of additions and exploits both node-level and feature-level parallelism. We also observe that the accelerator implements two modes of operation depending on the order of the matrix multiplications, which leads to different strategies for pipelining. To accommodate them, the modules are interconnected both from the aggregate module to the combination modules and vice versa.

Rubik. Similarly to the case above, Rubik [23] proposes a hardware accelerator assisted by some pre-processing in software. On the hardware side, Rubik presents a hierarchical PE array design, wherein each PE contains a number of MAC units plus instruction and data queues to feed them. The design is unified because aggregations and combinations are scheduled across all PEs. Moreover, each PE includes two small private caches that store recently accessed vertices and partial aggregations. Each PE is connected to the rest of PEs and two memory controllers placed on the side via a meshed NoC. On the software side, Rubik proposes lightweight graph reordering (once per graph) to put together nodes that are connected with each other, similarly to Reference [155], but here to improve the performance of the private PE caches.

GCNAX. The work in Reference [96] points out the load imbalance, execution order, and loop optimization inefficiencies from other accelerators, whose impact varies across workloads. To address them, the authors propose GCNAX as a flexible accelerator whose dataflow is reconfigurable in terms of loop order and loop fusion strategy. To find the most effective dataflows for each particular dataset, the authors perform a design space exploration of dataflow design decisions. Therefore, in inference, GCNAX is reconfigured based on the characteristics of the problem at hand. Finally, GCNAX uses the outer product to mitigate the effect of unbalanced presence of zeros, unlike other accelerators. Thanks to these techniques, GCNAX is around 10× and 2× faster and more efficient than HyGCN and AWB-GCN, respectively.

GraphACT. While all other accelerators focused on inference, GraphACT [175] explores how to efficiently perform GNN training in an heterogeneous CPU+FPGA platform. The main design decision relates to determining which parts are computed where and which data to store in memory. To address these questions, the authors argue that CPU performs graph sampling and the calculation of the loss gradients, while the FPGA does forward and backward propagation passes. The FPGA, thus implements aggregation and combination. The authors present optimizations based on the scheduling of the different operations taking into consideration that backpropagation can be performed after batching of multiple layers or batching different parts of the graph. Moreover,

similarly to in Reference [177], redundant operations at aggregation are eliminated via searching of edges common to multiple vertices.

4.3 Discussion

The analysis of the state of the art performed in previous sections leads to several conclusions. First, we observe that a quantitative comparison among systems is very difficult due to the lack of a common baseline system and a GNN benchmark suite with a representative set of algorithms, datasets, and design targets. To bridge this gap, initiatives such as the **Open Graph Benchmark (OGB)** [72] or GNNmark [10] aim to provide a representative set of graphs and GNNs to use as benchmarks. In hardware accelerators, comparing multiple recent architectures is difficult and some works have compared their fundamental dataflows instead [96]. In this direction, Garg et al. performed a dataflow classification that includes multiple operation orders, and whose analysis that may guide further developments in the field [51].

A second reflection is that the desirable *one approach fits all* does not apply to GNNs, and distinct design approaches will probably be required for different applications. For example, the extreme scale and high throughput demands of recommendation systems is well in line with the targets of software frameworks: programmability and scalability. In contrast, for applications that need to focus on real-time operation and energy efficiency, custom hardware acceleration solutions may be the only way to go. Moreover, the wide variety of problems with their different graph and feature vector sizes renders the acceleration problem more difficult to tackle with a single approach [103, 163, 182].

Finally, we identify a few outstanding challenges for acceleration. Support for dynamic graphs is a pending issue only evaluated in AliGraph [187]. Learning over dynamic graphs implies not only processing the GNN in each timestep but also updating the weight matrices as the graph evolves, factors that might be amenable to software or hardware optimization. At the frontier of software and hardware, another challenge resides in how to approach the GNN acceleration problem with a co-design strategy, i.e., which tasks can be offloaded to software and which ones should stay in hardware, taking into consideration the related overheads. On the hardware side, how to best accelerate training remains as an important open question as all proposals except GraphACT [175] have targeted inference. Beyond that, another challenge in hardware accelerators is finding the right balance between performance and generalization in light of the multitude of graph types and GNN variants, including techniques such as pooling, sampling, or skip connections.

5 GNN ACCELERATION: THE VISION

Previous sections have discussed how GNNs can be understood as a set of classical NNs working symbiotically over graph-structured data. We have seen that, to extract specific knowledge from the graphs, different NN layers may be employed leading to a wide variety of GNN flavours. This, plus the fundamental dependence of GNNs on the input graph (which may be extremely large) complicate the task of streamlining their execution. As a result, works on GNN acceleration have implicitly made a choice upon either providing an extremely efficient acceleration scheme for a specific GNN variant, or being general or flexible enough to serve multiple types of GNNs less efficiently.

The key challenge in GNN acceleration is thus to provide a framework that is able to both maximize performance and efficiency while maintaining a degree of flexibility that caters to the different graph sizes, characteristics, and GNN algorithms. Albeit a daunting task, in this section we aim to leverage the analysis of existing acceleration works to hypothesize that would be the main characteristics that future GNN accelerators should feature. In particular, our envisaged architectural approach shall be driven by (i) software-hardware co-design, (ii) graph awareness, and

Fig. 9. Architectural vision for GNN accelerators with hardware-software co-design (i.e., control and data planes), graph awareness (i.e., guided mapping and scheduling), and communication-centric design (i.e., re-configurable interconnect).

(iii) an much-needed emphasis on communications. We next discuss these aspects qualitatively, using Figure 9 as reference.

5.1 Software-Hardware Co-Design

The analysis of prior work has shown that both software and hardware approaches can provide significant speedups. In some occasions, one might argue that both strategies attack the problem similarly, e.g., node reordering in software [141] and workload balancing in hardware [53]. However, a few works have also started to realize that both approaches are not mutually exclusive and that their benefits can add up, or one can simplify the other. For instance, Rubik improves performance by reordering the graph in software [23]. Also, the design from Zhang et al. [177] eliminates redundant operations via software pre-processing and then optimizes execution with specialized aggregation and combination modules. The software side allows to avoid having specialized hardware structures to eliminate redundant operations.

Building upon this observation, our first proposed pillar is *software-hardware co-design* as a strategy for handling different GNNs and graphs efficiently while retaining some hardware simplicity. We advocate for a control-data plane model where, in general, the control plane will be implemented entirely in software providing the flexibility and the data plane will be implemented in custom hardware providing the efficiency. While conceptually separated (see Figure 9), the operation of both planes will be tightly coupled.

On the one hand, the *control plane* manages the actions of the accelerator by having a global view of the complete GNN structure and input graph. The control plane is responsible for dictating the dataflow running in the data plane, by (i) partitioning the GNN computation into manageable computational segments, (ii) mapping the different vertices and edges to the hardware resources of the data plane, and (iii) scheduling the different executions toward balancing the workload, maximize the benefits of pipelining, and so on. Finally, we also consider part of the control plane to (iv) drive pre-processing (and possibly offline) steps such as the removal of redundant operations [77] or the detection of certain graph aspects such as cliques [155]. By being implemented in software, all these functions can deliver the required flexibility to accelerate any GNN workload. However, given that certain pre-processing steps may take minutes or hours in very large graphs [28], care must be taken in not turning the software side into the bottleneck of the system. To this

end, one may resort to lightweight heuristics or limit software techniques to specific cases such as deep GNNs or training, where the result of pre-processing may be reused multiple times.

However, the *data plane* consists of the processing and memory elements that work as per the control plane instructions to execute a GNN. As we have seen in Section 4.2, we could adopt many strategies for architecting the data plane, e.g., unified, phased, modular, homogenous, heterogeneous, to name a few. However, we find particularly interesting the use of architectures similar to that of MAERI [90], where an homogeneous array of PEs and a specialized memory hierarchy are put together via a lightweight reconfigurable interconnect fabric. This architecture could adapt the dataflow according to the control plane commands, thus allowing to give service to the multiple execution stages of an algorithm or different algorithms.

5.2 Graph Awareness

Most accelerators have attempted to provide methods that adapt to runtime conditions while being largely unaware of the input graph characteristics [76, 103]. However, it has been also realized that aspects such as the size of the graph, the relative size of the feature vectors, the clustering factor of the graph, or the average degree of the same can be extremely relevant in accelerating the GNN [45, 141, 164]. In fact, GNNAdvisor [155] seeks to exploit this information explicitly to improve the performance in GPUs, while others have based the order of operations or the mapping of PEs on characteristics of the graph [53, 103, 164]. Other characterization works have shown that the impact of loop ordering or dataflow design decisions on performance certainly depends on the input graph [10, 51, 96].

This leads to the second pillar of our envisaged architecture: *graph awareness*. If the GNN depends on the input graph, then maximizing performance needs to be aware of the main features of that graph. Offline or online methods shall be used to extract useful information from the graph that, in our case, will be leveraged by the control plane. This will thus affect aspects such as the graph partitioning [141], which may be more or less aggressive depending on the degree distribution; the ordering and pipelining of the different aggregate–combine phases, which may vary across layers and across graphs; or the scheduling process to minimize inter-partition communication. A good example of this approach is community detection, whose efficient implementation [46, 110] or prediction [26] may allow for the partition of the graph in densely connected graphlets at runtime. This is relevant to efficient pooling [168], redundancy elimination [77], and optimal scheduling [155]. Again, it is critical to minimize the overhead of techniques providing graph awareness, either via heuristics, reuse of prior analyses, or its use only in certain occasions where the pre-processing can be done in advance or its benefit maximized, i.e., training.

5.3 Communication-Centric Design

Data movement is the enemy of efficient architectures. Hardware accelerators aim to minimize it by adapting its resources to the execution dataflow but, surprisingly, traditional DNN accelerators [39, 138] have generally given a relatively low importance to the sub-system handling data movement: the interconnect fabric. This is also true for GNN accelerators, which are generally computing-centric with few exceptions [144, 175]. However, GNNs pose the additional challenge of not having a single optimal dataflow given the input graph dependence and the many algorithm variants. Thus, data movement continues to be a crucial aspect [60].

For this reason, the third pillar of our envisaged architecture is taking a *communication-centric design* approach. This is a philosophy that has been applied to endow DNN accelerators with certain flexibility [88–90] or to optimize distributed learning [139]. In our case, we propose the use of a reconfigurable interconnect fabric among the PEs to adapt the hardware to the underlying graph connectivity or, in other words, to the optimal dataflow that may vary across layers, partitions, or

graphs. In an extreme case, one could adopt the approach of recent DNN accelerators that orchestrate all data movement at compilation time [4, 75]. GNNs and their extreme size might discourage the use of this strategy and, instead, advocate for a compilation that provides hints for the interconnect to adapt to the varying needs of the graph and its most optimal dataflow. The compilation and reconfiguration could be complemented by the analysis of the input graph. Assuming it can be done in advance or with little overhead, graph profiling may allow us to predict the prevalent communication patterns and, thus, the most appropriate interconnect topology.

6 CONCLUSION

The recent interest in geometric deep learning, or methods able to model and predict graph-structured data, have led to an explosion of research around GNNs. As we have seen in our analysis of the current state of the art, most of the works focus on the algorithms and their applications, rendering the topic of GNN computing a less beaten path. However, we anticipate that the area of software and hardware support for GNNs will grow at a fast pace, continuing an upwards trend that we observed from 2018 to today.

The reasons for the probable increase in research delving into more efficient computing means for GNNs are several. First, the field is maturing and the more theoretical algorithm-driven research gives way to the most application-oriented development. A clear example of this trend is the advent of efforts to unify aspects such as benchmarking [72]. Second, GNNs are the key to many disruptive applications in multiple fields, thus creating a clear application pull driving the need for better processing. Third, GNNs present multiple unique challenges such as the wide variety of algorithm variants, their dependence on the graph characteristics, or their massive scale in some applications. This makes the field of GNN processing unlikely to saturate in the foreseeable future and calls for an in-depth discussion of not only the challenges associated to GNN processing, but also of possible ways to tackle them.

Finally, we highlight the rising popularity of software frameworks and the recent appearance of hardware accelerators for GNNs. On the software side, libraries such as DGL or NeuGraph aim to speed up and add features to widespread frameworks such as TF or PyTorch. Interesting contributions are acceleration of GNNs via graph analysis or pre-coding, as well as the distribution of computation in large-scale systems, much needed for huge recommendation systems. On the hardware side, we did not observe a clear architectural trend and existing proposals are debating between being specific or applicable to multiple GNN variants, and between unified architectures or more hierarchical, tiled organizations. Building on this observation, we envision that future accelerators shall adopt a hardware-software co-design approach to maximize performance, keep graph awareness as a profitable optimization opportunity, and tackle workload variability via a reconfigurable interconnect.

ACKNOWLEDGMENTS

The authors thank the anonymous reviewers and the editorial team for their constructive criticism, which has helped improve the quality of the article. We are also grateful to Albert Cabellos-Aparicio, Tushar Krishna, José L. Abellán, and Manuel E. Acacio for the countless discussions on the topic.

REFERENCES

- [1] 2019. Build Graph Nets in Tensorflow. Retrieved from https://github.com/deepmind/graph_nets.
- [2] 2019. Food Discovery with Uber Eats: Using Graph Learning to Power Recommendations. Retrieved from <https://eng.uber.com/uber-eats-graph-learning>.
- [3] 2020. PaddlePaddle/PGL. Retrieved from <https://github.com/PaddlePaddle/PGL>.

- [4] Dennis Abts, Jonathan Ross, Jonathan Sparling, Mark Wong-VanHaren, Max Baker, Tom Hawkins, Andrew Bell, John Thompson, Temesghen Kahsai, Garrin Kimmell, et al. 2020. Think fast: A tensor streaming processor (TSP) for accelerating deep learning workloads. In *Proceedings of the ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA'20)*. 145–158.
- [5] Luis B. Almeida. 1990. A learning rule for asynchronous perceptrons with feedback in a combinatorial environment. In *Artificial Neural Networks: Concept Learning*. 102–111.
- [6] James Atwood and Don Towsley. 2016. Diffusion-convolutional neural networks. In *Advances in Neural Information Processing Systems*. 1993–2001.
- [7] Adam Auten, Matthew Tomei, and Rakesh Kumar. 2020. Hardware acceleration of graph neural networks. *Proceedings of the Design Automation Conference*.
- [8] Matej Balog, Bart van Merriënboer, Subhdeep Moitra, Yujia Li, and Daniel Tarlow. 2019. Fast training of sparse graph neural networks on dense hardware. arXiv:1906.11786. Retrieved from <https://arxiv.org/abs/1906.11786>.
- [9] Niccolò Bandinelli, Monica Bianchini, and Franco Scarselli. 2010. Learning long-term dependencies using layered graph neural networks. In *Proceedings of the International Joint Conference on Neural Networks*.
- [10] Trinayan Baruah, Kaustubh Shivdikar, Shi Dong, Yifan Sun, Saiful A. Mojumder, Kihoon Jung, José L. Abellán, Yash Ukidave, Ajay Joshi, John Kim, and David Kaeli. 2021. GNNMark: A benchmark suite to characterize graph neural network training on GPUs. In *Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS'21)*. IEEE, 13–23.
- [11] Peter W. Battaglia, Jessica B. Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, et al. 2018. Relational inductive biases, deep learning, and graph networks. arXiv:1806.01261. Retrieved from <https://arxiv.org/abs/1806.01261>.
- [12] Peter W. Battaglia, Razvan Pascanu, Matthew Lai, Danilo Rezende, and Koray Kavukcuoglu. 2016. Interaction networks for learning about objects, relations and physics. In *Advances in Neural Information Processing Systems*, 4502–4510.
- [13] Monica Bianchini, Marco Maggini, Lorenzo Sarti, and Franco Scarselli. 2005. Recursive neural networks for processing graphs with labelled edges: Theory and applications. *Neural Netw.* 18, 8 (2005), 1040–1050.
- [14] Matthew Botvinick, Sam Ritter, Jane X. Wang, Zeb Kurth-Nelson, Charles Blundell, and Demis Hassabis. 2019. Reinforcement learning, fast and slow. *Trends Cogn. Sci.* 23, 5 (2019), 408–422.
- [15] Xavier Bresson and Thomas Laurent. 2017. Residual gated graph convnets. arXiv:1711.07553. Retrieved from <https://arxiv.org/abs/1711.07553>.
- [16] Michael M. Bronstein, Joan Bruna, Yann Lecun, Arthur Szlam, and Pierre Vandergheynst. 2017. Geometric deep learning: Going beyond euclidean data. *IEEE Sign. Process. Mag.* 34, 4 (2017), 18–42.
- [17] Ivan Brugere, Brian Gallagher, and Tanya Y. Berger-Wolf. 2018. Network structure inference, a survey: Motivations, methods, and applications. *ACM Comput. Surv.* 51, 2 (2018).
- [18] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. 2014. Spectral networks and deep locally connected networks on graphs. *Proceedings of the 2nd International Conference on Learning Representations (ICLR'14)*.
- [19] Ines Chami, Sami Abu-El-Haija, Bryan Perozzi, Christopher Ré, and Kevin Murphy. 2020. Machine learning on graphs: A model and comprehensive taxonomy. arXiv:2005.03675. Retrieved from <https://arxiv.org/abs/2005.03675>.
- [20] Jie Chen, Tengfei Ma, and Cao Xiao. 2018. FastGCN: Fast learning with graph convolutional networks via importance sampling. In *Proceedings of the 6th International Conference on Learning Representations (ICLR'18)*.
- [21] Jianfei Chen, Jun Zhu, and Le Song. 2018. Stochastic training of graph convolutional networks with variance reduction. In *Proceedings of the 35th International Conference on Machine Learning (ICML'18)*, 1503–1532.
- [22] Ming Chen, Zhewei Wei, Zengfeng Huang, Bolin Ding, and Yaliang Li. 2020. Simple and deep graph convolutional networks. In *Proceedings of the International Conference on Machine Learning*. 1725–1735.
- [23] Xiaobing Chen, Yuke Wang, Xinfeng Xie, Xing Hu, Abanti Basak, Ling Liang, Mingyu Yan, Lei Deng, Yufei Ding, Zidong Du, and Yuan Xie. 2021. Rubik: A hierarchical architecture for efficient graph neural network training. *IEEE Trans. Comput.-Aided Des. Integr. Circ. Syst.* (2021). <https://ieeexplore.ieee.org/abstract/document/9428002>.
- [24] Y. Chen, T. Krishna, J. S. Emer, and V. Sze. 2017. Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks. *IEEE J. Solid-State Circ.* 52, 1 (2017), 127–138.
- [25] Y. Chen, T. Yang, J. Emer, and V. Sze. 2019. Eyeriss v2: A flexible accelerator for emerging deep neural networks on mobile devices. *IEEE J. Emerg. Select. Top. Circ. Syst.* 9, 2 (2019), 292–308.
- [26] Zhengdao Chen, Joan Bruna, and Lisha Li. 2019. Supervised community detection with line graph neural networks. *Proceedings of the 7th International Conference on Learning Representations (ICLR'19)*.
- [27] Zhiqian Chen, Fanglan Chen, Lei Zhang, Taoran Ji, Kaiqun Fu, Liang Zhao, Feng Chen, and Chang-Tien Lu. 2020. Bridging the gap between spatial and spectral domains: A survey on graph neural networks. arXiv:2002.11867. Retrieved from <https://arxiv.org/abs/2002.11867>.

- [28] Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. 2019. Cluster-gcn: An efficient algorithm for training deep and large graph convolutional networks. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 257–266.
- [29] Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. 2019. Cluster-GCN: An efficient algorithm for training deep and large graph convolutional networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- [30] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. In *Proceedings of the 8th Workshop on Syntax, Semantics and Structure in Statistical Translation (SSST-8'14)*.
- [31] Peng Cui, Xiao Wang, Jian Pei, and Wenwu Zhu. 2019. A survey on network embedding. *IEEE Trans. Knowl. Data Eng.* 31, 5 (2019), 833–852.
- [32] Zhiyong Cui, Kristian Henrickson, Ruimin Ke, and Yin Hai Wang. 2020. Traffic graph convolutional recurrent neural network: A deep learning framework for network-scale traffic learning and forecasting. *IEEE Trans. Intell. Transport. Syst.* 21, 11 (2020), 4883–4894.
- [33] Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, Luke Vilnis, Ishan Durugkar, Akshay Krishnamurthy, Alex Smola, and Andrew McCallum. 2018. Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning. In *Proceedings of the 7th International Conference on Learning Representations*.
- [34] Shail Dave, Riyadh Baghdadi, Tony Nowatzki, Sasikanth Avancha, Aviral Shrivastava, and Baoxin Li. 2020. Hardware acceleration of sparse and irregular tensor computations of ML models: A survey and insights. arXiv:2007.00864. <https://arxiv.org/abs/2007.00864>.
- [35] Nicola De Cao and Thomas Kipf. 2018. MolGAN: An implicit generative model for small molecular graphs. In *Proceedings of the ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models*.
- [36] Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: Simplified data processing on large clusters. *Commun. ACM* 51, 1 (2008), 107–113.
- [37] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*. 3844–3852.
- [38] Vincenzo Di Massa, Gabriele Monfardini, Lorenzo Sarti, Franco Scarselli, Marco Maggini, and Marco Gori. 2006. A comparison between recursive neural networks and graph neural networks. *Proceedings of the IEEE International Conference on Neural Networks*, 778–785.
- [39] Zidong Du, Robert Fashhuber, Tianshi Chen, Paolo Ienne, Ling Li, Tao Luo, Xiaobing Feng, Yunji Chen, and Olivier Temam. 2015. ShiDianNao: Shifting vision processing closer to the sensor. In *Proceedings of the ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA'15)*. 92–104.
- [40] Alberto Garcia Duran and Mathias Niepert. 2017. Learning graph representations with embedding propagation. In *Advances in Neural Information Processing Systems*. 5119–5130.
- [41] David Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P. Adams. 2015. Convolutional networks on graphs for learning molecular fingerprints. *Adv. Neural Inf. Process. Syst.* (2015), 2224–2232.
- [42] Vijay Prakash Dwivedi, Chaitanya K Joshi, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. 2020. Benchmarking graph neural networks. In *Proceedings of the ICML Workshop on Graph Representation Learning and Beyond*.
- [43] Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542, 7639 (2017), 115–118.
- [44] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. *Proceedings of the The Web Conference (WWW'19)*, 417–426.
- [45] Matthias Fey and Jan Eric Lenssen. 2019. Fast graph representation learning with PyTorch Geometric. *Proceedings of the International Conference on Learning Representations (ICLR'19)*.
- [46] Santo Fortunato. 2010. Community detection in graphs. *Phys. Rep.* 486, 3 (2010), 75–174.
- [47] Alex Fout, Jonathon Byrd, Basir Shariat, and Asa Ben-Hur. 2017. Protein interface prediction using graph convolutional networks. *Adv. Neural Inf. Process. Syst.* 6531–6540.
- [48] Hongyang Gao, Zhengyang Wang, and Shuiwang Ji. 2018. Large-scale learnable graph convolutional networks. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2018).
- [49] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. 2020. VectorNet: Encoding HD maps and agent dynamics from vectorized representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'20)*. 11522–11530.
- [50] Victor Garcia and Joan Bruna. 2018. Few-shot learning with graph neural networks. *Proceedings of the International Conference on Learning Representations*.
- [51] Raveesh Garg, Eric Qin, Francisco Muñoz Martínez, Robert Guirado, Akshay Jain, Sergi Abadal, José L. Abellán, Manuel E. Acacio, Eduard Alarcón, Sivasankaran Rajamanickam, and Tushar Krishna. 2021. A taxonomy for

- classification and comparison of dataflows for GNN accelerators. arXiv:2103.07977. Retrieved from <https://arxiv.org/abs/2103.07977>.
- [52] Thomas Gärtner, Peter Flach, and Stefan Wrobel. 2003. On graph kernels: Hardness results and efficient alternatives. In *Learning Theory and Kernel Machines*, Bernhard Schölkopf and Manfred K. Warmuth (Eds.). Springer, Berlin, 129–143.
 - [53] Tong Geng, Ang Li, Runbin Shi, Chunshu Wu, Tianqi Wang, Yanfei Li, Pouya Haghi, Antonino Tumeo, Shuai Che, Steve Reinhardt, and Martin C. Herbordt. 2020. AWB-GCN: A graph convolutional network accelerator with runtime workload rebalancing. In *Proceedings of the 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO'20)*. IEEE, 922–936.
 - [54] Xu Geng, Yaguang Li, Leye Wang, Lingyu Zhang, Qiang Yang, Jieping Ye, and Yan Liu. 2019. Spatiotemporal multi-graph convolution network for ride-hailing demand forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*. 3656–3663.
 - [55] Swarnendu Ghosh, Nibaran Das, Teresa Gonçalves, and Paulo Quaresma. 2018. The journey of graph kernels through two decades. *Comput. Sci. Rev.* 27 (2018), 88–111.
 - [56] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. 2017. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning*. 2053–2070.
 - [57] Marco Gori, Gabriele Monfardini, and Franco Scarselli. 2005. A new model for learning in Graph domains. In *Proceedings of the International Joint Conference on Neural Networks*. 729–734.
 - [58] Daniele Grattarola and Cesare Alippi. 2021. Graph neural networks in tensorflow and keras with spektral [Application Notes]. *IEEE Comput. Intell. Mag.* 16, 1 (2021), 99–106.
 - [59] Chuang-Yi Gui, Long Zheng, Bingsheng He, Cheng Liu, Xin-Yu Chen, Xiao-Fei Liao, and Hai Jin. 2019. A survey on graph processing accelerators: Challenges and opportunities. *J. Comput. Sci. Technol.* 34, 2 (2019), 339–371.
 - [60] Robert Guirado, Akshay Jain, Sergi Abadal, and Eduard Alarcón. 2021. Characterizing the communication requirements of GNN accelerators: A model-based approach. In *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS'21)*.
 - [61] Shengnan Guo, Youfang Lin, Ning Feng, Chao Song, and Huaiyu Wan. 2019. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 922–929.
 - [62] Erkam Guresen, Gulgun Kayakutlu, and Tugrul U. Daim. 2011. Using artificial neural network models in stock market index prediction. *Expert Syst. Appl.* 38, 8 (2011), 10389–10397.
 - [63] Tae Jun Ham, Lisa Wu, Narayanan Sundaram, Nadathur Satish, and Margaret Martonosi. 2016. Graphicionado: A high-performance and energy-efficient accelerator for graph analytics. In *Proceedings of the Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-49)*. IEEE.
 - [64] Takuo Hamaguchi, Hidekazu Oiwai, Masashi Shimbo, and Yuji Matsumoto. 2017. Knowledge transfer for out-of-knowledge-base entities: A graph neural network approach. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 1802–1808.
 - [65] William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*. 1025–1035.
 - [66] William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Representation learning on graphs: Methods and applications. *IEEE Data Eng. Bull.* 40, 3 (2017), 52–74.
 - [67] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
 - [68] Kartik Hegde, Jiyong Yu, Rohit Agrawal, Mengjia Yan, Michael Pellauer, and Christopher Fletcher. 2018. UCNN: Exploiting computational reuse in deep neural networks via weight repetition. In *Proceedings of the ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA'18)*. 674–687.
 - [69] Mikael Henaff, Joan Bruna, and Yann LeCun. 2015. Deep convolutional networks on graph-structured data. arXiv:1506.05163. Retrieved from <https://arxiv.org/abs/1506.05163>.
 - [70] Tamás Horváth, Thomas Gärtner, and Stefan Wrobel. 2004. Cyclic pattern kernels for predictive graph mining. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 158–167.
 - [71] Fenyu Hu, Yanqiao Zhu, Shu Wu, Liang Wang, and Tieniu Tan. 2019. Hierarchical graph convolutional networks for semi-supervised node classification. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
 - [72] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. 2020. Open graph benchmark: Datasets for machine learning on graphs. In *Advances in Neural Information Processing Systems*, Vol. 33. 22118–22133.
 - [73] Yuwei Hu, Zihao Ye, Minjie Wang, Jiali Yu, Da Zheng, Mu Li, Zheng Zhang, Zhiru Zhang, and Yida Wang. 2020. FeatGraph: A flexible and efficient backend for graph neural network systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*.

- [74] Guyue Huang, Guohao Dai, Yu Wang, and Huazhong Yang. 2020. GE-SpMM: General-purpose Sparse Matrix-Matrix Multiplication on GPUs for Graph Neural Networks. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. Article 72.
- [75] Michael James, Marvin Tom, Patrick Groeneveld, and Vladimir Kibardin. 2020. Physical mapping of neural networks on a wafer-scale deep learning accelerator. In *Proceedings of the International Symposium on Physical Design*. 145–149.
- [76] Zhihao Jia, Sina Lin, Mingyu Gao, Matei Zaharia, and Alex Aiken. 2020. Improving the accuracy, scalability, and performance of Graph Neural Networks with ROC. In *Proceedings of the Conference on Machine Learning and Systems (MLSys'20)*.
- [77] Zhihao Jia, Sina Lin, Rex Ying, and Alex Aiken. 2020. Redundancy-Free computation for graph neural networks. In *Proceedings of the SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'20)*.
- [78] Zhihao Jia, Matei Zaharia, and Alex Aiken. 2019. Beyond data and model parallelism for deep neural networks. *Proceedings of the Conference on Machine Learning and Systems (SysML'19)*.
- [79] Xiaodong Jiang, Pengsheng Ji, and Sheng Li. 2019. CensNet: Convolution with edge-node switching in graph neural networks. *Proceedings of the International Joint Conference on Artificial Intelligence (2019)*, 2656–2662.
- [80] Xiangyang Ju, Steven Farrell, Paolo Calafiura, Daniel Murnane, Prabhat, Lindsey Gray, et al. 2019. Graph neural networks for particle reconstruction in high energy physics detectors. In *Proceedings of the 2nd Workshop on Machine Learning and the Physical Sciences (NeurIPS'19)*.
- [81] George Karypis and Vipin Kumar. 1998. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.* 20, 1 (1998), 359–392.
- [82] Tatsuro Kawamoto, Masashi Tsubaki, and Tomoyuki Obuchi. 2018. Mean-field theory of graph neural networks in graph partitioning. In *Advances in Neural Information Processing Systems*. 4361–4371.
- [83] Byung-Hoon Kim and Jong Chul Ye. 2020. Understanding graph isomorphism network for rs-fMRI functional connectivity analysis. *Front. Neurosci.* 14 (2020), 630.
- [84] Kevin Kinningham, Philip Levis, and Christopher Re. 2020. GrETA: Hardware optimized graph processing for GNNs. In *Proceedings of the Workshop on Resource-Constrained Machine Learning (ReCoML'20)*.
- [85] Kevin Kinningham, Christopher Re, and Philip Levis. 2020. GRIP: A graph neural network accelerator architecture. arXiv:2007.13828. Retrieved from <https://arxiv.org/abs/2007.13828>.
- [86] Thomas N. Kipf and Max Welling. 2016. Variational graph auto-encoders. In *Proceedings of the Bayesian Deep Learning Workshop (NIPS'16)*.
- [87] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. *Proceedings of the 5th International Conference on Learning Representations*.
- [88] Hyoukjun Kwon and Tushar Krishna. 2017. Rethinking NoCs for spatial neural network accelerators. In *Proceedings of the IEEE/ACM International Symposium on Networks-on-Chip (NOCS'17)*.
- [89] Hyoukjun Kwon, Ananda Samajdar, and Tushar Krishna. 2018. A communication-centric approach for designing flexible DNN accelerators. *IEEE Micro* 38, 6 (2018), 25–35.
- [90] Hyoukjun Kwon, Ananda Samajdar, and Tushar Krishna. 2018. MAERI: Enabling flexible dataflow mapping over DNN accelerators via reconfigurable interconnects. In *Proceedings of the 23rd International Conference on Architectural Support for Programming Languages and Operating Systems*. 461–475.
- [91] Luis C. Lamb, Artur d'Ávila Garcez, Marco Gori, Marcelo O. R. Prates, Pedro H. C. Avelar, and Moshe Y. Vardi. 2020. Graph neural networks meet neural-symbolic computing: A survey and perspective. (2020), 4877–4884. <https://arxiv.org/pdf/2003.00330.pdf>.
- [92] Yann Lecun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (2015), 436–444.
- [93] John Boaz Lee, Ryan A. Rossi, Sungchul Kim, Nesreen K. Ahmed, and Eunye Koh. 2019. Attention models in graphs: A survey. *ACM Trans. Knowl. Discov. Data* 13, 6, Article 62 (2019), 1–25.
- [94] Adam Lerer, Ledell Wu, Jiajun Shen, Timothee Lacroix, Luca Wehrstedt, Abhijit Bose, and Alex Peysakhovich. 2019. PyTorch-BigGraph: A large-scale graph embedding system. In *Proceedings of the Conference on Machine Learning and Systems (MLSys'19)*.
- [95] Guohao Li, Matthias Muller, Ali Thabet, and Bernard Ghanem. 2019. DeepGCNs: Can GCNs go as deep as CNNs?. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9267–9276.
- [96] Jiajun Li, Ahmed Louri, Avinash Karanth, and Razvan Bunescu. 2021. GCNAX: A flexible and energy-efficient accelerator for graph convolutional neural networks. In *Proceedings of the IEEE International Symposium on High-Performance Computer Architecture (HPCA'21)*. IEEE, 775–788.
- [97] Qimai Li, Zhichao Han, and Xiao Ming Wu. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*. 3538–3545.
- [98] Ruiyu Li, Makarand Tapaswi, Renjie Liao, Jiaya Jia, Raquel Urtasun, and Sanja Fidler. 2017. Situation recognition with graph neural networks. In *Proceedings of the IEEE International Conference on Computer Vision* 4183–4192.

- [99] Ruoyu Li, Sheng Wang, Feiyun Zhu, and Junzhou Huang. 2018. Adaptive graph convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [100] Xiaoxiao Li, Yuan Zhou, Nicha C. Dvornek, Muhan Zhang, Juntang Zhuang, Pamela Ventola, and James S. Duncan. 2020. Pooling regularized graph neural network for fMRI biomarker analysis. In *Proceedings of the Medical Image Computing and Computer Assisted Intervention (MICCAI'20)*. 625–635.
- [101] Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. 2018. Learning deep generative models of graphs. In *Proceedings of the International Conference on Learning Representations (ICLR'18) Workshops*.
- [102] Yujia Li, Richard Zemel, Marc Brockschmidt, and Daniel Tarlow. 2016. Gated graph sequence neural networks. In *Proceedings of the 4th International Conference on Learning Representation*.
- [103] Shengwen Liang, Ying Wang, Cheng Liu, Lei He, Huawei Li, and Xiaowei Li. 2021. EnGN: A high-throughput and energy-efficient accelerator for large graph neural networks. *IEEE Trans. Comput.* 70, 9 (2021), 1511–1525.
- [104] Renjie Liao, Marc Brockschmidt, Daniel Tarlow, Alexander L. Gaunt, Raquel Urtasun, and Richard S. Zemel. 2018. Graph partition neural networks for semi-supervised classification. *Proceedings of the International Conference on Learning Representations (ICLR'18) Workshops*.
- [105] Husong Liu, Shengliang Lu, Xinyu Chen, and Bingsheng He. 2020. G3: When graph neural networks meet parallel graph processing systems on GPUs. *VLDB Endow.* 13, 12 (2020), 2813–2816.
- [106] Lingxiao Ma, Zhi Yang, Youshan Miao, Jilong Xue, Ming Wu, Lidong Zhou, and Yafei Dai. 2019. Neugraph: Parallel deep neural network computation on large graphs. In *Proceedings of the USENIX Annual Technical Conference (USENIX ATC'19)*. 443–458.
- [107] Tianle Ma and Aidong Zhang. 2019. AffinityNet: Semi-supervised few-shot learning for disease type prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 1069–1076.
- [108] Yuzhe Ma, Haoxing Ren, Brucek Khailany, Harbinder Sikka, Lijuan Luo, Karthikeyan Natarajan, and Bei Yu. 2019. High performance graph convolutional networks with applications in testability analysis. *Proceedings of the Design Automation Conference (2019)*.
- [109] Yanjun Ma, Dianhai Yu, Tian Wu, and Haifeng Wang. 2019. PaddlePaddle: An open-source deep learning platform from industrial practice. *Front. Data Comput.* 1, 1 (2019), 105–115.
- [110] Fragkiskos D. Malliaros and Michalis Vazirgiannis. 2013. Clustering and community detection in directed networks: A survey. *Phys. Rep.* 533, 4 (2013), 95–142.
- [111] Sparsh Mittal. 2020. A survey of FPGA-based accelerators for convolutional neural networks. *Neural Comput. Appl.* 32 (2020), 1109–1139.
- [112] Gabriele Monfardini, Vincenzo Di Massa, Franco Scarselli, and Marco Gori. 2006. Graph neural networks for object localization. *Front. Artif. Intell. Appl.* 141 (2006), 665–669.
- [113] Federico Monti, Michael M. Bronstein, and Xavier Bresson. 2017. Geometric matrix completion with recurrent multi-graph neural networks. In *Advances in Neural Information Processing Systems*. 3698–3708.
- [114] Vojtech Myska, Radim Burget, and Brezany Peter. 2019. Graph neural network for website element detection. *Proceedings of the 42nd International Conference on Telecommunications and Signal Processing*, 216–219.
- [115] Yusuke Nagasaka, Akira Nukada, Ryosuke Kojima, and Satoshi Matsuoka. 2019. Batched sparse matrix multiplication for accelerating graph convolutional networks. In *Proceedings of the 19th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. 231–240.
- [116] David F. Nettleton. 2013. Data mining of social networks represented as graphs. *Comput. Sci. Rev.* 7 (2013), 1–34.
- [117] Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. 2015. A review of relational machine learning for knowledge graphs. *Proc. IEEE* 104, 1 (2015), 11–33.
- [118] Mathias Niepert, Mohamed Ahmad, and Konstantin Kutzkov. 2016. Learning convolutional neural networks for graphs. *Proceedings of the 33rd International Conference on Machine Learning*. 2958–2967.
- [119] Daniel Oñoro-Rubio, Mathias Niepert, Alberto García-Durán, Roberto González, and Roberto J. López-Sastre. 2017. Answering visual-relational queries in web-extracted knowledge graphs. In *Proceedings of the Annual Conference on Automated Knowledge Base Construction (AKBC'17)*.
- [120] Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao Schardl, and Charles Leiserson. 2020. EvolveGCN: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 5363–5370.
- [121] Hogan Park and Jennifer Neville. 2019. Exploiting interaction links for node classification with deep graph neural networks. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- [122] Fernando J. Pineda. 1987. Generalization of back-propagation to recurrent neural networks. *Phys. Rev. Lett.* 59, 19 (1987), 2229–2232.
- [123] Jiezhong Qiu, Jian Tang, Hao Ma, Yuxiao Dong, Kuansan Wang, and Jie Tang. 2018. DeepInf: Social influence prediction with deep learning. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (2018)*.

- [124] Afshin Rahimi, Trevor Cohn, and Timothy Baldwin. 2018. Semi-supervised user geolocation via graph convolutional networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*. 2009–2019.
- [125] Krzysztof Rusek and Piotr Cholda. 2019. Message-passing neural networks learn little’s law. *IEEE Commun. Lett.* 23, 2 (2019), 274–277.
- [126] Krzysztof Rusek, José Suárez-Varela, Albert Mestres, Pere Barlet-Ros, and Albert Cabellos-Aparicio. 2019. Unveiling the potential of Graph Neural Networks for network modeling and optimization in SDN. In *Proceedings of the ACM Symposium on SDN Research* (2019).
- [127] Guillaume Salha, Romain Hennequin, and Michalis Vazirgiannis. 2019. Keep it simple: Graph autoencoders without graph convolutional networks. In *Proceedings of the NeurIPS’19 Graph Representation Learning Workshop*.
- [128] Alvaro Sanchez-Gonzalez, Nicolas Heess, Jost Tobias Springenberg, Josh Merel, Martin Riedmiller, Raia Hadsell, and Peter Battaglia. 2018. Graph networks as learnable physics engines for inference and control. In *Proceedings of the 35th International Conference on Machine Learning*. 7097–7117.
- [129] Benjamin Sanchez-Lengeling, Jennifer N. Wei, Brian K. Lee, Richard C. Gerkin, Alán Aspuru-Guzik, and Alexander B. Wiltschko. 2019. Machine learning for scent: Learning generalizable perceptual representations of small molecules. arXiv:1910.10685. Retrieved from <https://arxiv.org/abs/1910.10685>.
- [130] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and G. Monfardini. 2009. The Graph Neural Network Model. *IEEE Trans. Neural Netw.* 20, 1 (2009), 61–80.
- [131] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2009. Computational capabilities of graph neural networks. *IEEE Trans. Neural Netw.* 20, 1 (2009), 81–102.
- [132] Franco Scarselli, Markus Hagenbuchner, Sweah Liang Yong, Ah Chung Tsoi, Marco Gori, and Marco Maggini. 2005. Graph neural networks for ranking web pages. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence*. 666–672.
- [133] Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. Modeling Relational Data with Graph Convolutional Networks. *Lecture Notes in Computer Science* (2018), 593–607.
- [134] Weijing Shi and Raj Rajkumar. 2020. Point-gnn: Graph neural network for 3d object detection in a point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1711–1719.
- [135] Hy Truong Son and Chris Jones. 2019. Graph neural networks with efficient tensor operations in CUDA/GPU and Graphflow deep learning framework in C++ for quantum chemistry. Retrieved from <http://people.cs.uchicago.edu/~hytruongson/CCN-GraphFlow.pdf>.
- [136] Peter C. St. John, Caleb Phillips, Travis W. Kemper, A. Nolan Wilson, Yanfei Guan, Michael F. Crowley, Mark R. Nimlos, and Ross E. Larsen. 2019. Message-passing neural networks for high-throughput polymer screening. *J. Chem. Phys.* 150, 23 (2019).
- [137] Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. 2016. Learning multiagent communication with backpropagation. In *Advances in Neural Information Processing Systems*. 2244–2252.
- [138] Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel S. Emer. 2017. Efficient processing of deep neural networks: A tutorial and survey. *Proc. IEEE* 105, 12 (2017), 2295–2329.
- [139] Zhenheng Tang, Shaohuai Shi, Xiaowen Chu, Wei Wang, and Bo Li. 2020. Communication-efficient distributed deep learning: A comprehensive survey. arXiv:2003.06307v1. Retrieved from <https://arxiv.org/abs/2003.06307v1>.
- [140] Kiran K. Thekumparampil, Chong Wang, Sewoong Oh, and Li-Jia Li. 2018. Attention-based graph neural network for semi-supervised learning. arXiv:1803.03735. Retrieved from <https://arxiv.org/abs/1803.03735>.
- [141] Chao Tian, Lingxiao Ma, Zhi Yang, and Yafei Dai. 2020. PCGCN: Partition-centric processing for accelerating graph convolutional network. *Proceedings of the IEEE International Parallel & Distributed Processing Symposium (IPDPS’20)* (2020), 936–945.
- [142] Dominika Tkaczyk, Paweł Szostek, Mateusz Fedoryszak, Piotr Jan Dendek, and Łukasz Bolikowski. 2015. CERMINE: Automatic extraction of structured metadata from scientific literature. *Int. J. Doc. Anal. Recogn.* 18, 4 (2015), 317–335.
- [143] Dinh V. Tran, Nicolò Navarin, and Alessandro Sperduti. 2018. On filter size in graph convolutional networks. In *Proceedings of the IEEE Symposium Series on Computational Intelligence (SSCI’18)*. IEEE, 1534–1541.
- [144] Alok Tripathy, Katherine Yelick, and Aydin Buluc. [n.d.]. Reducing communication in graph neural network training. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. 987–1000.
- [145] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 6000–6010.
- [146] Petar Veličković, Arantxa Casanova, Pietro Liò, Guillem Cucurull, Adriana Romero, and Yoshua Bengio. 2018. Graph attention networks. In *Proceedings of the 6th International Conference on Learning Representations*.
- [147] Manisha Verma and Debasis Ganguly. 2019. Graph edit distance computation via graph neural networks yunsheng. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1281–1284.

- [148] Saurabh Verma and Zhi Li Zhang. 2019. Stability and generalization of graph convolutional neural networks. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1539–1548.
- [149] Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu. 2016. CNN-RNN: A unified framework for multi-label image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2285–2294.
- [150] Lei Wang, Qiang Yin, Chao Tian, Jianbang Yang, Rong Chen, Wenyuan Yu, Zihang Yao, and Jingren Zhou. 2021. FlexGraph: A flexible and efficient distributed framework for GNN training. In *Proceedings of the 16th European Conference on Computer Systems*. 67–82.
- [151] Minjie Wang, Lingfan Yu, Da Zheng, Quan Gan, Yu Gai, Zihao Ye, Mufei Li, Jinjing Zhou, Qi Huang, Chao Ma, et al. 2019. Deep graph library: Towards efficient and scalable deep learning on graphs. arXiv:1909.01315. Retrieved from <https://arxiv.org/abs/1909.01315>.
- [152] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. 2018. Non-local Neural Networks. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 7794–7803.
- [153] Xuhong Wang, Ding Lyu, Mengjian Li, Yang Xia, Qi Yang, Xinwen Wang, Xinguang Wang, Ping Cui, Yupu Yang, Bowen Sun, et al. 2021. APAN: Asynchronous propagation attention network for real-time temporal graph embedding. In *Proceedings of the International Conference on Management of Data*. 2628–2638.
- [154] Yangzihao Wang, Andrew Davidson, Yuechao Pan, Yuduo Wu, Andy Riffel, and John D Owens. 2016. Gunrock: A high-performance graph processing library on the GPU. In *Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*.
- [155] Yuke Wang, Boyuan Feng, Gushu Li, Shuangchen Li, Lei Deng, Yuan Xie, and Yufei Ding. 2021. GNNAdvisor: An efficient runtime system for GNN acceleration on GPUs. In *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation (OSDI'21)*.
- [156] Jim Webber. 2012. A programmatic introduction to neo4j. In *Proceedings of the 3rd Annual Conference on Systems, Programming, and Applications: Software for Humanity*.
- [157] Boris Weisfeiler and Andrei Leman. 1968. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Tekhn. Inf.* (1968), 2–16.
- [158] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. 2019. Simplifying graph convolutional networks. In *Proceedings of the International Conference on Machine Learning*. 6861–6871.
- [159] Jun Wu, Jingrui He, and Jiejun Xu. 2019. Demo-Net: Degree-specific graph neural networks for node and graph classification. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 406–415.
- [160] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2021. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* 32, 1 (2021), 4–24.
- [161] Zhipu Xie, Weifeng Lv, Shangfo Huang, Zhilong Lu, and Bowen Du. 2020. Sequential graph neural network for urban road traffic speed prediction. *IEEE Access* 8 (2020), 63349–63358.
- [162] Keyulu Xu, Stefanie Jegelka, Weihua Hu, and Jure Leskovec. 2019. How powerful are graph neural networks? *Proceedings of the 7th International Conference on Learning Representations* (2019).
- [163] Mingyu Yan, Zhaodong Chen, Lei Deng, Xiaochun Ye, Zhimin Zhang, Dongrui Fan, and Yuan Xie. 2020. Characterizing and understanding GCNs on GPU. *IEEE Comput. Arch. Lett.* 19, 1 (2020), 22–25.
- [164] Mingyu Yan, Lei Deng, Xing Hu, Ling Liang, Yujing Feng, Xiaochun Ye, Zhimin Zhang, Dongrui Fan, and Yuan Xie. 2020. HyGCN: A GCN accelerator with hybrid architecture. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA'20)*. 15–29.
- [165] Sijie Yan, Yuanjun Xiong, and Dahua Lin. 2018. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [166] Yiding Yang, Xinchao Wang, Mingli Song, Junsong Yuan, and Dacheng Tao. 2019. SPAGAN: Shortest Path Graph Attention Network. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- [167] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L. Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 974–983.
- [168] Rex Ying, Jiaxuan You, Christopher Morris, Xiang Ren, William L. Hamilton, and Jure Leskovec. 2018. Hierarchical graph representation learning with differentiable pooling. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. 4805–4815.
- [169] Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. 2019. Gnnexplainer: Generating explanations for graph neural networks. In *Advances in Neural Information Processing Systems*. 9244–9255.
- [170] Sweah Liang Yong, Markus Hagenbuchner, Ah Chung Tsoi, Franco Scarselli, and Marco Gori. 2006. Document mining using graph neural network. In *Proceedings of the International Workshop of the Initiative for the Evaluation of XML Retrieval*. Springer, 458–472.

- [171] T. Young, D. Hazarika, S. Poria, and E. Cambria. 2018. Recent trends in deep learning based natural language processing. *IEEE Comput. Intell. Mag.* 13, 3 (2018), 55–75.
- [172] Bing Yu, Haoteng Yin, and Zhanxing Zhu. 2018. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*.
- [173] Wenchao Yu, Cheng Zheng, Wei Cheng, Charu C Aggarwal, Dongjin Song, Bo Zong, Haifeng Chen, and Wei Wang. 2018. Learning deep network representations with adversarially regularized autoencoders. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- [174] Victoria Zayats and Mari Ostendorf. 2018. Conversation modeling on Reddit using a graph-structured LSTM. *Trans. Assoc. Comput. Ling.* 6 (2018), 121–132.
- [175] Hanqing Zeng and Viktor Prasanna. 2020. GraphACT: Accelerating GCN training on CPU-FPGA heterogeneous platforms. In *Proceedings of the ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*. 255–265.
- [176] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2019. GraphSAINT: Graph sampling based inductive learning method. In *Proceedings of the International Conference on Learning Representations*.
- [177] Bingyi Zhang, Hanqing Zeng, and Viktor Prasanna. 2020. Hardware acceleration of large scale GCN inference. In *Proceedings of the IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP'20)*. 61–68.
- [178] Dalong Zhang, Xin Huang, Ziqi Liu, Jun Zhou, Zhiyang Hu, Xianzheng Song, Zhibang Ge, Lin Wang, Zhiqiang Zhang, and Yuan Qi. 2020. AGL: A scalable system for industrial-purpose graph machine learning. *VLDB Endow.* 13, 12 (2020), 3125–3137.
- [179] Li Zhang, Dan Xu, Anurag Arnab, and Philip HS Torr. 2020. Dynamic graph message passing networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3726–3735.
- [180] Muhan Zhang and Yixin Chen. 2018. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems*.
- [181] Ziwei Zhang, Peng Cui, and Wenwu Zhu. 2020. Deep learning on graphs: A survey. *IEEE Trans. Knowl. Data Eng.* 14, 8 (2020).
- [182] Zhihui Zhang, Jingwen Leng, Lingxiao Ma, Youshan Miao, Chao Li, and Minyi Guo. 2020. Architectural implications of graph neural networks. *IEEE Comput. Arch. Lett.* 19, 1 (2020), 59–62.
- [183] Ziwei Zhang, Chenhao Niu, Peng Cui, Bo Zhang, Wei Cui, and Wenwu Zhu. 2020. A simple and general graph neural network with stochastic message passing. arXiv:2009.02562. Retrieved from <https://arxiv.org/abs/2009.02562>.
- [184] Da Zheng, Chao Ma, Minjie Wang, Jinjing Zhou, Qidong Su, Xiang Song, Quan Gan, Zheng Zhang, and George Karypis. 2020. DistDGL: Distributed graph neural network training for billion-scale graphs. In *Proceedings of the IEEE/ACM Workshop on Irregular Applications: Architectures and Algorithms (IA3'20)*.
- [185] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2020. Graph neural networks: A review of methods and applications. *AI Open* 1 (2020), 57–81.
- [186] Di Zhu and Yu Liu. 2018. Modelling spatial patterns using graph convolutional networks. *Leibniz Int. Proc. Inf.* 114 (2018).
- [187] Rong Zhu, Kun Zhao, Hongxia Yang, Wei Lin, Chang Zhou, Baole Ai, Yong Li, and Jingren Zhou. 2018. AliGraph: A comprehensive graph neural network platform. *VLDB Endow.* (2018), 2094–2105.
- [188] Daniel Zügner and Stephan Günnemann. 2019. Adversarial attacks on graph neural networks via meta learning. In *Proceedings of the 7th International Conference on Learning Representations (2019)*.

Received November 2020; revised July 2021; accepted July 2021