

Deep Learning for Community Detection: Progress, Challenges and Opportunities

Fanzhen Liu¹, Shan Xue^{1,2}, Jia Wu^{1*}, Chuan Zhou³, Wenbin Hu⁴,
Cecile Paris^{2,1}, Surya Nepal^{2,1}, Jian Yang¹, Philip S. Yu⁵

¹ Department of Computing, Macquarie University, Sydney, Australia

² CSIRO's Data61, Sydney, Australia

³ Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, China

⁴ School of Computer Science, Wuhan University, Wuhan, China

⁵ Department of Computer Science, University of Illinois at Chicago, Chicago, USA
{fanzhen.liu@hdr., jia.wu@, jian.yang@}mq.edu.au, zhouchuan@amss.ac.cn, hwb@whu.edu.cn,
{emma.xue, surya.nepal, cecile.paris}@data61.csiro.au, psyu@uic.edu

Abstract

As communities represent similar opinions, similar functions, similar purposes, etc., community detection is an important and extremely useful tool in both scientific inquiry and data analytics. However, the classic methods of community detection, such as spectral clustering and statistical inference, are falling by the wayside as deep learning techniques demonstrate an increasing capacity to handle high-dimensional graph data with impressive performance. Thus, a survey of current progress in community detection through deep learning is timely. Structured into three broad research streams in this domain – deep neural networks, deep graph embedding, and graph neural networks, this article summarizes the contributions of the various frameworks, models, and algorithms in each stream along with the current challenges that remain unsolved and the future research opportunities yet to be explored.

1 Introduction

The two basic elements of a network/graph are nodes and edges. From the perspective of connectedness and density, communities are known as locally dense connected subgraphs or clusters of nodes [Fortunato and Hric, 2016]. In addition to the internal cohesion of subgraphs, their separation from each other should also be taken into account. To this end, graph theory sets out two specific rules for determining which of these nodes and edges to form a community: 1) nodes in a community are densely connected; and 2) nodes in different communities are sparsely connected. An easy and popular understanding is that communities are subgraphs holding more internal connections than external connections.

In the real world, the nodes in a community can share common properties or serve similar functions [Fortunato, 2010], and finding these commonalities is the strategic key to almost every community detection strategy that exists today. Community detection, or more specifically, clustering nodes based

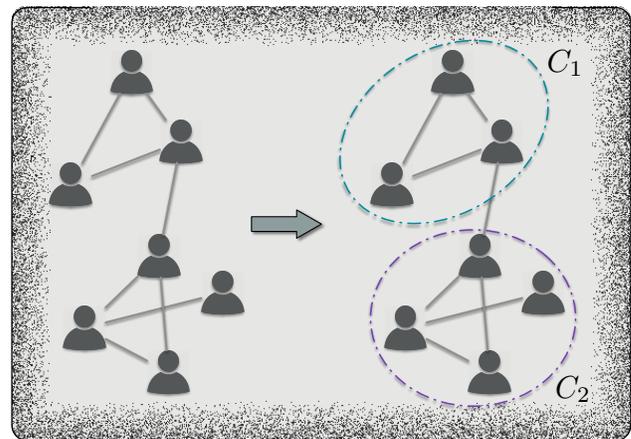


Figure 1: An example of community detection in a social network. The network is divided into two communities according to the closeness of the individuals: community C_1 includes three nodes, and community C_2 includes four nodes.

on a similar feature or set of features, helps us understand the inherent patterns and functions of networks. For example, community detection in protein-protein interaction (PPI) networks reveals proteins with similar biological functions [Chen and Yuan, 2006]. In citation networks, community detection determines the importance, interconnectedness, and evolution of research topics [Chen and Redner, 2010]. In enterprise networks, employees can be grouped into communities by studying company's offline internal sources and online enterprise social relationships [Zhang *et al.*, 2017]. In social networks like Twitter and Facebook, users with common interests or mutual friends may be members of the same community [Yang *et al.*, 2013], as shown in Figure 1.

Most traditional methods of community detection are based on statistical inference and conventional machine learning. For example, one of the most widely-used methods to detect communities and describe how they are formed is stochastic block model [Karrer and Newman, 2011]. However, despite its past popularity, this strategy struggles in the

*Corresponding Author

face of today's complex datasets and complicated social scenarios. In conventional machine learning, detecting communities has generally been conceived as a clustering problem on graphs. But these approaches are highly dependent on the characteristics of the data. For example, spectral clustering [Ng *et al.*, 2002], which uses eigenvectors to partition nodes into communities, does not perform well with sparse networks.

It is becoming ever-clearer that the increasing scale of networks and dimensionality of data demand more powerful techniques to maintain effective and efficient performance with feasible computation speeds [Rosvall and Bergstrom, 2008].

For now at least, deep learning [LeCun *et al.*, 2015] is the solution. With deep learning, computational models can learn representations of data at multiple levels of abstraction, which is perfectly suited to network data. Moreover, its ability to learn nonlinear features is greatly advancing, and it has been successful in a wide range of fields where the data have internal relationships, such as computer vision and natural language processing. Further, multi-layer deep neural networks can reduce the dimensionality of the data [Zhang, 2018], which widens the potential scope of network analysis tasks like community detection, node classification, and link prediction.

In this survey, we focus on new trends in deep learning for community detection. Our investigation is broken into three main parts:

- 1) We review and assess the advantages of the various deep learning methods for community detection;
- 2) We summarize and categorize the state-of-the-art studies from a technical perspective; and
- 3) We identify and discuss the technical challenges yet to be solved, along with suggestions of promising opportunities for future works.

To the best of our knowledge, this is the first work to provide a comprehensive review of deep learning for community detection. Most of the papers we surveyed were published in recent influential international conferences in the domains of artificial intelligence, machine learning, data mining, and web services, *e.g.*, NIPS, ICLR, AAAI, IJCAI, KDD, ICDM, CIKM, ICDE, and WWW. Articles in high-quality peer-reviewed journals were also included.

The information in this survey should directly support researchers and technology experts to understand the past, current, and future trends in fields of deep learning for community detection.

The remainder of this survey is organized as follows. Section 2 clarifies the concepts of community detection. Section 3 explains the reasons why today's community detection methods must be based on deep learning. Novel deep learning-based community detection methods are surveyed in Section 4. Section 5 discusses the remaining challenges in this field given recent progress, along with suggestions for future research opportunities for academics and domain experts. Finally, Section 6 concludes the survey.

2 Community Detection

In this paper, a network is a special graph that abstracts complex relationships in real-world systems, such as the internets, academic collaborations, or social groups.

Definition 1 (Network) *Following graph theory, a weighted network is represented as $G = (V, E, W)$ and an unweighted network is represented as $G = (V, E)$, where V and E denote the set of nodes and edges, respectively, and W denotes the corresponding weights of E in terms of the intensity or capacity of the connections. In an unweighted network, W is regarded as 1 and can be removed from G .*

A subgraph $g \subseteq G$ is a partition of a graph that retains the original network structure. The division of subgraphs follows pre-defined rules, and different rules may result in different forms of subgraph. A community is a type of subgraph that represents a real social phenomenon. In other words, a community is a group of people or objects that share common characteristics.

Definition 2 (Community) *Communities are the subgraphs in a network where the nodes share dense connections. Sparsely-connected nodes delineate communities. Here, we use $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$ to denote a set of k communities divided from a network G , where C_i is the i th community from this network partition. A node v clustered into the community C_i satisfies the condition that the internal degree of each node inside the community exceeds its external degree.*

Thus, the goal of community detection is to discover the communities \mathcal{C} in a network G .

3 Why Detection by Deep Learning?

The clear advantage of deep learning for community detection over other machine learning methods is its ability to encode feature representations of high-dimensional data [Zhang, 2018]. With graph-structured data, that translates to leveraging the network's topological information [Fortunato, 2010]. Deep learning models can also learn the pattern of nodes, neighborhoods, and subgraphs [Wu *et al.*, 2020]. Plus, they are much more resilient to the sparsity associated with large-scale graph data. Lastly, in many real-world scenarios, considering that the majority of nodes are unlabeled and there is little to no prior knowledge of the communities within the data, deep learning is the superior choice for unsupervised learning tasks [Tian *et al.*, 2014].

Beyond simply examining network topologies for detecting communities, some strategies also explore semantic descriptions as node features in the data. Traditional community detection approaches are mainly based on an adjacency matrix and a node attribute matrix [Yang *et al.*, 2013; He *et al.*, 2017]. Deep learning, however, can create much more powerful representations of node attributes and community structures. To this end, recent studies have indicated promising new directions in community detection with deep learning – for example, the deep learning models in [Xie *et al.*, 2018; Bhatia and Rani, 2019; Sun *et al.*, 2019] and modified deep learning models based on community properties [Zhang *et al.*, 2018; Tu *et al.*, 2019].

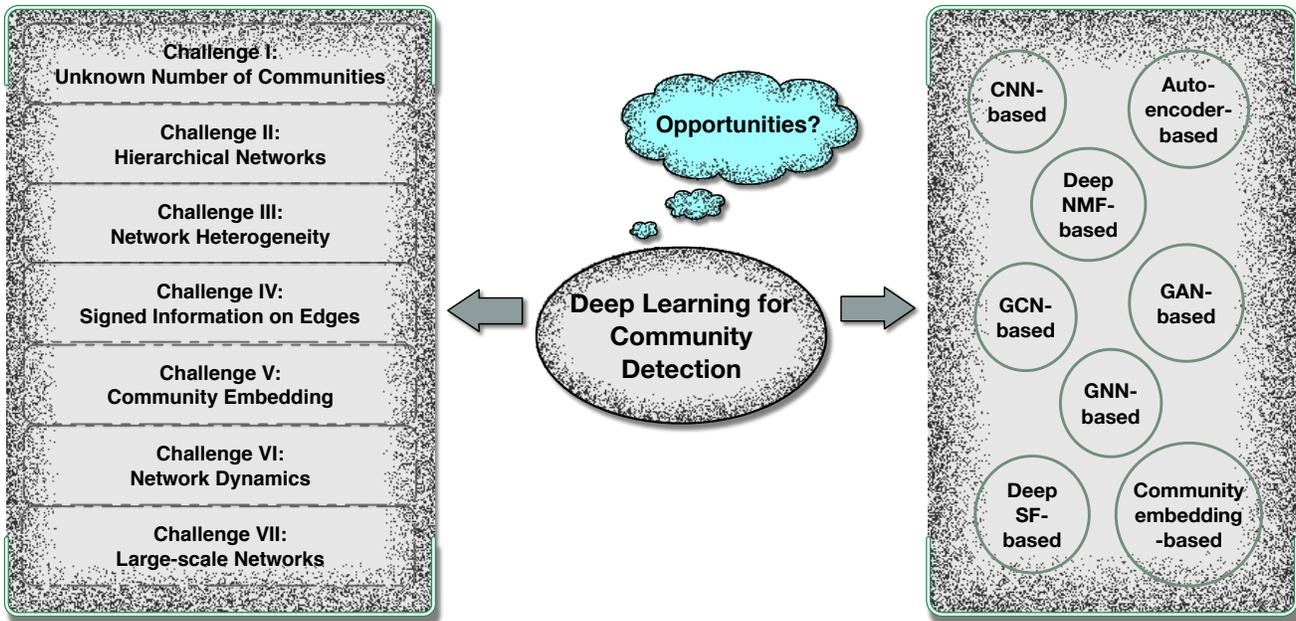


Figure 2: Deep learning for community detection: progress, challenges, and opportunities. This overview is based on concepts from the community detection literature on convolutional neural networks (CNNs), auto-encoders, generative adversarial networks (GANs), deep non-negative matrix factorization (NMF), deep sparse filtering (SF), community embedding, graph neural networks (GNNs), and graph convolutional networks (GCNs). Each approach is discussed in more detail in Section 4. Each challenge and related opportunity is discussed in more detail in Section 5.

These studies provide us a glimpse of what deep learning might bring to the future of community detection:

- performance improvements;
- the capacity to base detection on more and richer features;
- the capacity to base detection on both network topology and node attributes for a more robust and better-performing model; and
- the ability to detect more complex structures in larger-scale networks.

4 Community Detection with Deep Learning

This section provides a technical overview of the recent research progress in deep learning for community detection. Each subsection covers one of the three broad categories of approaches, being deep neural networks, deep graph embedding, and graph neural networks. Within each subsection, we summarize the most influential methods and the contributions of the various frameworks, models, and algorithms. An overview of these methods and their corresponding challenges is provided in Figure 2.

4.1 Deep Neural Network-based Community Detection

Deep neural networks have a natural strength for modeling and capturing comprehensive relationships. The three most

popular deep neural network models in the community detection domain are convolutional neural networks (CNNs), auto-encoders, and generative adversarial networks (GANs). Each is summarized in turn below.

CNN-based Approaches

The two key components of CNNs are the convolution operation and the pooling operation over the convolutional layers. The convolution operation uses convolution kernels to reduce computation costs. The pooling operation is subsequently applied to the feature mapping to ensure the robustness of CNNs.

Leveraging advancements in CNNs, [Xin *et al.*, 2017] designed a novel CNN model to detect communities in topologically incomplete networks, where some edges are missing when observed from real-world networks. [Sperlí, 2019] included sparse matrix convolution within a CNN framework specifically to deal with the highly sparse representations associated with adjacency matrices.

Auto-encoder-based Approaches

Stacked auto-encoders are a very powerful form of deep neural network model for community detection [Yang *et al.*, 2016].

The inspiration to apply auto-encoders to community detection came from the discovery that auto-encoders and spectral clustering have similar frameworks in terms of a low-dimensional approximation of the spectral matrix [Cao *et al.*, 2018b]. Focusing on network topology, the approach devised by [Bhatia and Rani, 2018] learns the node clusters via

a random walk-based personalized PageRank and fine-tunes the detection by optimizing the modularity of the community structure. To utilize the node attribute information, [Cao *et al.*, 2018b] proposed a stacked auto-encoder that combines community detection via both the network topology and the attributes of the nodes to enhance the generalization ability of the hidden layer of the deep neural network. To further address the matches between the network topology and node attributes, [Cao *et al.*, 2018a] developed a graph regularized auto-encoder approach by introducing an adaptive parameter as a trade-off control for the matches.

To avoid the need to preset the number of communities, a layer-wise stacked auto-encoder can effectively find centers of communities based on the network structure [Bhatia and Rani, 2018]. Further, this automatic selection mechanism ensures the model assigns nodes strictly based on community criteria. Since this discovery, more and more approaches adaptively learn community structures instead of pre-defining a number. For example, [Choong *et al.*, 2018] introduced a Mixture of Gaussian to capture higher-order patterns from community structures and model the generative process of the observed network to detect communities.

For networks where connections are with positive and negative signs, they are named as signed networks [Xu *et al.*, 2019]. To handle the signed information on edges, a semi-supervised stacked auto-encoder reconstructs the adjacency matrix to represent the signed network so as to the further learning of deep network embedding [Shen and Chung, 2018].

GAN-based Approaches

Generative adversarial networks (GANs) involve two deep neural networks competing with each other, which results in fast-adjusting training precision. These methods typically run unsupervised, generating new data with the same statistics as the training set. GANs have been explored to graph representation tasks with great effectiveness [Wang *et al.*, 2019].

[Jia *et al.*, 2019] argued that traditional graph clustering-based community detection methods cannot handle the dense overlapping of communities, where one node may belong to more than one community. To this end, they proposed a novel model CommunityGAN that jointly solves overlapping community detection and graph representation learning based on GANs. More importantly, compared with the general graph node representations that have no specific meanings, CommunityGAN has the capability to represent the membership strength of nodes to communities.

4.2 Deep Graph Embedding-based Community Detection

Deep graph embedding is a technique that maps nodes in the network to a low-dimensional vector space, while saving as much structural information as possible in the representations [Xue *et al.*, 2019]. Graph embedding approaches are particularly suited to machine learning tasks based on network analysis, *e.g.*, link prediction, node classification, and node clustering, because they can reference latent features in the representation instead of searching the network. After that, the

clustering methods, such as k -means, can support the community detection.

Deep NMF-based Approaches

The non-negative matrix factorization (NMF) is a group of computational algorithms that factorize a matrix into two matrices with the property that all three matrices have no negative elements. For community detection, NMF approaches approximate the adjacency matrix of a network into the product of two factorized matrices by minimizing the error function for the further clustering task.

[Ye *et al.*, 2018] proposed a novel deep NMF model that boosts performance by mapping the community structures back to the original network structure during deep learning. By adding node attributes and forming an attributed graph, deep NMF-based community detection only requires an adjacency matrix and node attribute matrix for factorization. In addition, [Li *et al.*, 2018] used NMF for attributed community detection based on node attributes and community embedding in line with the deep feature learning and deep network embedding.

Deep SF-based Approaches

Embeddings can encode the input of pairwise relations to avoid searching a sparse adjacency matrix, and hence sparse filtering (SF) is an effective deep feature learning algorithm that only requires one hyperparameter to process high-dimensional inputs [Ngiam *et al.*, 2011]. [Xie *et al.*, 2018] proposed an efficient network representation approach for network community discovery by way of deep sparse filtering that works particularly well with large-scale networks. An unsupervised deep learning algorithm extracts the network features, which are then used to partition the network.

Community Embedding-based Approaches

Traditionally, graph embedding focuses on individual nodes, but communities in networks reflect high-order proximities, such as similar opinions and behaviors. This is an important yet largely under-explored area of graph embedding focused on embedding communities, where the goal is to learn the node distributions of communities in a low-dimensional space. [Cavallari *et al.*, 2017] investigated this idea, arguing that this novel and non-trivial strategy could be beneficial to community detection. For example, through a transitional graph embedding method, node distributions could be used to preserve the network structure so as to improve community detection in reverse. [Zhang *et al.*, 2018] proposed a community-preserving network embedding method to learn network representations. Their performance on community detection demonstrates the superiority. In addition, [Tu *et al.*, 2019] proposed a novel graph embedding model that learns the embeddings of both the nodes and the communities, and its optimization process alternates between community assignment and node embedding, instead of simultaneously solving both tasks.

4.3 Graph Neural Network-based Community Detection

GNNs are a technical fusion of graph mining and deep learning, and their recent rapid development is an indication of

their power to model and capture the complex relationships in graph-based data. For example, GNN for supervised community detection in [Chen *et al.*, 2019] introduced a non-backtracking operator to define the edge adjacency. Not only is this approach a feasible way to improve learning performance, but the operator selection is also convenient.

Graph Convolutional Networks (GCNs) inherit the fast learning of CNNs and further improve on that benefit by integrating a probabilistic model that considers the probability distribution of the entities in the network. For example, [Jin *et al.*, 2019] combined a Markov random field with an attributed network of semantic information to support semi-supervised learning, while [Shchur and Gunnemann, 2019] integrated a Bernoulli–Poisson probabilistic model with a GCN for overlapping community detection such that the convolutional layers can recognize complex network patterns.

5 Challenges and Opportunities

Although we have witnessed the rapid developments in deep learning for community detection in recent years, especially the last five, there are still issues that need better solutions and challenges that remain unresolved. As with most gaps in our knowledge, these problems provide opportunities for future study. In this section, we discuss seven broad challenges facing the community, beginning with, arguably, the longest-standing issue in community detection.

5.1 Challenge I: An Unknown Number of Communities

The challenge caused by the need to know the number of communities in advance has existed for a very long time, and, despite some professed solutions, this problem still has not been fully resolved. In machine learning, community detection is usually approached as an unsupervised clustering task because most of the data extracted from real-world networks do not have labels. This leads to a catch-22: the data needs to be labeled to determine the number of communities, but the number of communities needs to be known before the data can be labeled. Deep learning methods work around this problem to some extent by clustering nodes according to similarity in one or more latent spaces. However, the clustering algorithms still need to know the final number of clusters in advance.

Opportunities: To this point, an effective solution is to calculate the number of communities by analyzing the network topology, such as [Bhatia and Rani, 2018; Bhatia and Rani, 2019] based on random walk-based personalized PageRank. However, this type of method cannot guarantee that every node in the network is assigned to a community. Thus, a full and complete solution to this problem is yet to be found.

5.2 Challenge II: Hierarchical Networks

Hierarchical networks are made up of layered networks, where each layer of networks shares specified functions. Hence, the community detection strategy must be able to extract layer-wise representations. The challenges involved in this include distinguishing different relationship types, such as horizontal and vertical, and managing varying levels of sparsity in different layers.

Opportunities: [Song and Thiagarajan, 2018] proposed a multi-layer DeepWalk with creating inter-layer edges to exploit the dependencies across different layers while preserving hierarchical structures. It learns representations for each node in every layer which are fine-tuned by a refinement strategy. Another possible solution is to simultaneously optimize for common representations applicable to all layers and local representations preserving the layer-specific network structure. Moreover, the scalability of the scheme with respect to the number of layers in the hierarchy is questionable and should be considered in devising novel solutions. In addition, novel models are desired to distinguish different kinds of connections in hierarchical networks. Thus, there is much work to be done before we have a deep learning method for detecting hierarchical network communities.

5.3 Challenge III: Network Heterogeneity

Network heterogeneity refers to networks that contain significantly different types of entities and relationships, which means the strategies used for homogeneous networks do not necessarily work. In particular, the different probability distributions associated with each type of entity need to be addressed in the design of models and algorithms.

Opportunities: To date, few deep learning approaches consider network heterogeneity. Of the work we surveyed, [Chang *et al.*, 2015] tackled this issue with a nonlinear embedding function to capture the complex interactions between heterogeneous network entities. However, their method neglects the different semantics of relationships between nodes [Fu *et al.*, 2017]. The future opportunities for community detection in heterogeneous networks may include: 1) deep graph embedding models and supporting algorithms; 2) specific deep learning models with novel training processes to learn heterogeneous graph properties within the hidden layers; and 3) novel models that can exploit different types of relationships among nodes.

5.4 Challenge IV: Signed Information on Edges

Many real-world networks have signed edges representing positive or negative relationships. The challenge is that these relationships need to be treated in different ways.

Opportunities: One possible solution is to incorporate positive and negative edges by designing a random walk procedure. [Hu *et al.*, 2019] followed this idea and developed a sparse graph embedding model based on word embedding. But their method does not perform as well as the baseline spectral methods for some small real-world signed networks. Another possible solution is to reconstruct the adjacency matrix representation of a signed network. However, this raises other problems because, in the real-world, adjacency connections are overwhelmingly positive. [Shen and Chung, 2018] imposed a penalty to ensure their stacked auto-encoder model focuses more on reconstructing the rarer negative edges over the abundant positive edges. However, without prior knowledge of most of the relationships in a community, which is realistic in many situations, this approach will not work. Thus, efficient methods for unsupervised community detection in signed networks are still required.

5.5 Challenge V: Community Embedding

Community embedding looks to create representations of communities instead of representations for each individual node [Cavallari *et al.*, 2017]. Thus, focus shifts to community-aware high-order proximity instead of the 1- or 2-order proximity associated with node neighborhoods. This is an emerging research area with three main challenges to overcome: 1) high computational costs; 2) relational evaluations between node and community structures; and 3) other problems when applying deep learning models, such as distribution shifts across communities.

Opportunities: From our survey, we intuited an intelligent way to support community embedding by automatically selecting the representation modules on nodes and/or communities. To this end, we suggest the following research objectives: 1) explore how to integrate community embedding into a deep learning model; 2) determine how to directly embed community structures for a range of benefits, such as fast computation; and 3) devise a method for optimizing the hyperparameters in an integrated deep community detection learning model.

5.6 Challenge VI: Network Dynamics

Changing dynamics can affect either the network topology or the node attributes, each of which must be handled in their own way. Topological changes, such as adding or deleting a node or edge, not only cause changes in a local community, they can also lead to global changes across an entire network [Liu *et al.*, 2020]. With dynamic networks, deep learning models need to be re-trained over a series of snapshots. The technical challenge with the temporal attributes of a dynamic network lies in the deep feature extraction of dynamics.

Opportunities: Deep learning methods have not yet been well developed for detecting communities with dynamic spatial and temporal properties. Therefore, future directions of research include: 1) detecting and recognizing the spatial changes over communities; 2) learning deep patterns that embed both temporal feature and community structure information; and 3) developing a unified deep learning approach for community detection that can simultaneously handle both spatial and temporal features.

5.7 Challenge VII: Large-scale Networks

Today, large-scale networks can contain millions of nodes, edges, and structural patterns and can also be highly dynamic, as networks like Facebook and Twitter demonstrate. This presents many areas of development that need to be addressed before truly large-scale community detection can become a reality. For instance, large-scale networks may have their inherent scale characteristics, such as scale-free (*i.e.*, a power-law degree distribution) in social networks. Such particular distribution can influence the performance of deep learning in community detection. Scalability seems to be another key issue to enable deep learning to detect communities in large-scale networked environments. Constantly changing network typologies further increases the detection difficulty. Overall, deep community detection in large-scale networks involves all six of the preceding challenges as well as the challenge of scalable learning.

Opportunities: To fully use the rich information in large-scale networks, novel unsupervised clustering algorithms are needed that have lower computational complexity and greater flexibility. Distributed computing architectures are popular in large-scale machine learning community. Thus, one possible direction is to develop a robust deep learning detection approach that can achieve an high-performance collaborative computing. On the other hand, regarding the high-dimensional adjacency matrix, the key strategy of dimension reduction commonly used in deep learning (*i.e.*, matrix low-rank approximation) does not work with large-scale networks, and even the current distributed computing solutions are still too expensive. Therefore, novel deep learning frameworks, models, and algorithms are in high demand. As probably the biggest challenge in community detection, those frameworks need to far exceed the current benchmarks in precision and speed.

6 Conclusions

Nowadays, all of us live, work, and play in an intricate system of many communities. Likewise, as our knowledge of the world grows, we find objects perform more than one function, ideas serve more than one purpose, and there is far greater interconnectedness between all things than we previously imagined. Detecting these communities and their inherent functions and features helps us to comprehensively understand our surrounding environment and discover complex relationships to describe and explain social phenomena.

Traditional community detection methods have typically relied on statistical inference and conventional machine learning methods, such as stochastic block model and spectral clustering. But advancements in deep learning have given rise to community detection strategies with more powerful capacity of handling high-dimensional graph data with impressive performance.

In this survey, we reviewed the technical trends and current state-of-play in model and algorithm development in three broad deep-learning approaches to detecting communities in various scenarios. We also identified seven vital challenges that need to be overcome for community detection through deep learning to fully mature.

To some extent, these challenges serve as guidelines for the next generation of community detection. Although many researchers are actively engaged in addressing these problems, we point to the numerous opportunities that still exist to contribute to the development of this important ever-evolving field.

Acknowledgments

This work is supported in part by ARC DECRA under grant DE200100964, in part by the 2018 Collaborative Project between MQ and Data61, in part by the MQEPS under grant 96804590, in part by the NSFC under grants 61872360 and 61976162, and the Youth Innovation Promotion Association CAS under grant 2017210, and in part by NSF under grants III-1526499, III-1763325, III-1909323, and CNS-1930941. F. Liu and S. Xue contributed equally to this work.

References

- [Bhatia and Rani, 2018] V. Bhatia and R. Rani. DFuzzy: A deep learning-based fuzzy clustering model for large graphs. *Knowl. Inf. Syst.*, 57(1):159–181, 2018.
- [Bhatia and Rani, 2019] V. Bhatia and R. Rani. A distributed overlapping community detection model for large graphs using autoencoder. *Future Gener. Comput. Syst.*, 94:16–26, 2019.
- [Cao *et al.*, 2018a] J. Cao, D. Jin, and J. Dang. Autoencoder based community detection with adaptive integration of network topology and node contents. In *KSEM*, pages 184–196, 2018.
- [Cao *et al.*, 2018b] J. Cao, D. Jin, L. Yang, and J. Dang. Incorporating network structure with node contents for community detection on large networks using deep learning. *Neurocomputing*, 297:71–81, 2018.
- [Cavallari *et al.*, 2017] S. Cavallari, V. Zheng, H. Cai, K. Chang, and E. Cambria. Learning community embedding with community detection and node embedding on graphs. In *CIKM*, page 377–386, 2017.
- [Chang *et al.*, 2015] S. Chang, W. Han, J. Tang, G. Qi, C. Aggarwal, and T. Huang. Heterogeneous network embedding via deep architectures. In *KDD*, pages 119–128, 2015.
- [Chen and Redner, 2010] P. Chen and S. Redner. Community structure of the physical review citation network. *J. Informetr.*, 4(3):278–290, 2010.
- [Chen and Yuan, 2006] J. Chen and B. Yuan. Detecting functional modules in the yeast protein-protein interaction network. *Bioinformatics*, 22(18):2283–2290, 07 2006.
- [Chen *et al.*, 2019] Z. Chen, L. Li, and J. Bruna. Supervised community detection with line graph neural networks. In *ICLR*, 2019.
- [Choong *et al.*, 2018] J. Choong, X. Liu, and T. Murata. Learning community structure with variational autoencoder. In *ICDM*, pages 69–78, 2018.
- [Fortunato and Hric, 2016] S. Fortunato and D. Hric. Community detection in networks: A user guide. *Phys. Rep.*, 659:1–44, 2016. Community detection in networks: A user guide.
- [Fortunato, 2010] S. Fortunato. Community detection in graphs. *Phys. Rep.*, 486(3-5):75–174, 2010.
- [Fu *et al.*, 2017] T. Fu, W. Lee, and Z. Lei. Hin2vec: Explore meta-paths in heterogeneous information networks for representation learning. In *CIKM*, page 1797–1806, 2017.
- [He *et al.*, 2017] D. He, Z. Feng, D. Jin, X. Wang, and W. Zhang. Joint identification of network communities and semantics via integrative modeling of network topologies and node contents. In *AAAI*, pages 116–124, 2017.
- [Hu *et al.*, 2019] B. Hu, H. Wang, X. Yu, W. Yuan, and T. He. Sparse network embedding for community detection and sign prediction in signed social networks. *J. Ambient Intell. Hum. Comput.*, 10(1):175–186, Jan 2019.
- [Jia *et al.*, 2019] Y. Jia, Q. Zhang, W. Zhang, and X. Wang. CommunityGAN: Community detection with generative adversarial nets. In *WWW*, pages 784–794, 2019.
- [Jin *et al.*, 2019] D. Jin, Z. Liu, W. Li, D. He, and W. Zhang. Graph convolutional networks meet markov random fields: Semi-supervised community detection in attribute networks. In *AAAI*, pages 152–159, 2019.
- [Karrer and Newman, 2011] B. Karrer and M. Newman. Stochastic blockmodels and community structure in networks. *Phys. Rev. E*, 83:016107, 2011.
- [LeCun *et al.*, 2015] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521:436–444, 2015.
- [Li *et al.*, 2018] Y. Li, C. Sha, X. Huang, and Y. Zhang. Community detection in attributed graphs: an embedding approach. In *AAAI*, pages 338–345, 2018.
- [Liu *et al.*, 2020] F. Liu, J. Wu, S. Xue, C. Zhou, J. Yang, and Q. Sheng. Detecting the evolving community structure in dynamic social networks. *World Wide Web*, 23:715–733, 2020.
- [Ng *et al.*, 2002] A. Ng, M. Jordan, and Y. Weiss. On spectral clustering: analysis and an algorithm. In *NIPS*, pages 849–856, 2002.
- [Ngiam *et al.*, 2011] J. Ngiam, Z. Chen, S. Bhaskar, P. Koh, and A. Ng. Sparse filtering. In *NIPS*, pages 1125–1133, 2011.
- [Rosvall and Bergstrom, 2008] M. Rosvall and C. Bergstrom. Maps of random walks on complex networks reveal community structure. *PNAS*, 105(4):1118–1123, 2008.
- [Shchur and Gunnemann, 2019] O. Shchur and S. Gunnemann. Overlapping community detection with graph neural networks. In *KDD Workshop DLG’19*, 2019.
- [Shen and Chung, 2018] X. Shen and F. Chung. Deep network embedding for graph representation learning in signed networks. *IEEE Trans. Cybern.*, pages 1–8, 2018.
- [Song and Thiagarajan, 2018] H. Song and J. Thiagarajan. Improved deep embeddings for inferencing with multi-layered networks. *arXiv:1811.12156*, 2018.
- [Sperlí, 2019] G. Sperlí. A deep learning based community detection approach. In *SAC*, pages 1107–1110, 2019.
- [Sun *et al.*, 2019] F. Sun, M. Qu, J. Hoffmann, C. Huang, and J. Tang. vGraph: A generative model for joint community detection and node representation learning. In *NIPS*, pages 512–522, 2019.
- [Tian *et al.*, 2014] F. Tian, B. Gao, Q. Cui, E. Chen, and T. Liu. Learning deep representations for graph clustering. In *AAAI*, pages 1293–1299, 2014.
- [Tu *et al.*, 2019] C. Tu, X. Zeng, H. Wang, Z. Zhang, Z. Liu, M. Sun, B. Zhang, and L. Lin. A unified framework for community detection and network representation learning. *IEEE Trans. Knowl. Data Eng.*, 31(6):1051–1065, 2019.
- [Wang *et al.*, 2019] H. Wang, J. Wang, J. Wang, M. Zhao, W. Zhang, F. Zhang, W. Li, X. Xie, and M. Guo. Learning graph representation with generative adversarial nets. *IEEE Trans. Knowl. Data Eng.*, 2019.
- [Wu *et al.*, 2020] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Netw. Learn. Syst.*, 2020.
- [Xie *et al.*, 2018] Y. Xie, M. Gong, S. Wang, and B. Yu. Community discovery in networks with deep sparse filtering. *Pattern Recognit.*, 81:50–59, 2018.
- [Xin *et al.*, 2017] X. Xin, C. Wang, X. Ying, and B. Wang. Deep community detection in topologically incomplete networks. *Physica A*, 469:342–352, 2017.
- [Xu *et al.*, 2019] P. Xu, W. Hu, J. Wu, and B. Du. Link prediction with signed latent factors in signed social networks. In *KDD*, page 1046–1054, 2019.
- [Xue *et al.*, 2019] S. Xue, J. Lu, and G. Zhang. Cross-domain network representations. *Pattern Recognit.*, 94:135–148, 2019.
- [Yang *et al.*, 2013] J. Yang, J. McAuley, and J. Leskovec. Community detection in networks with node attributes. In *ICDM*, 2013.
- [Yang *et al.*, 2016] L. Yang, X. Cao, D. He, C. Wang, X. Wang, and W. Zhang. Modularity based community detection with deep learning. In *IJCAI*, pages 2252–2258, 2016.
- [Ye *et al.*, 2018] F. Ye, C. Chen, and Z. Zheng. Deep autoencoder-like nonnegative matrix factorization for community detection. In *CIKM*, pages 1393–1402, 2018.
- [Zhang *et al.*, 2017] J. Zhang, P. S. Yu, and Y. Lv. Enterprise community detection. In *ICDE*, pages 219–222, 2017.
- [Zhang *et al.*, 2018] Y. Zhang, T. Lyu, and Y. Zhang. Cosine: Community-preserving social network embedding from information diffusion cascades. In *AAAI*, pages 2620–2627, 2018.
- [Zhang, 2018] X. Zhang. Multilayer bootstrap networks. *Neural Networks*, 103:29–43, 2018.