Network Science Meets AI: A Converging Frontier

Matteo Zignani¹, Fragkiskos D. Malliaros², Ingo Scholtes³, Roberto Interdonato^{4,5}, Manuel Dileo¹

¹ Department of Computer Science, University of Milan, Italy
² Université Paris-Saclay, CentraleSupélec, Inria, France
³ CAIDAS, Julius-Maximilians-Universität Würzburg, Germany
⁴ CIRAD, UMR TETIS, Montpellier, France
⁵ INRIA, Montpellier, France

Abstract. The convergence of network science and artificial intelligence (AI) represents a rich area of research, where both fields can mutually enhance one another. Network science offers a comprehensive framework to analyze and model complex relationships, while machine learning (ML) and AI provide powerful tools for recognizing patterns and making predictions from large datasets. Combining these two disciplines can advance the study of complex systems and lead to new innovations in data-driven research. This tutorial paper reviews fundamental concepts of network science, describes the current and promising research direction for bridging network science and AI, and summarizes the contributions that have been accepted for publication in the ESANN 2025 special session on the topic.

1 Introduction

Network science and artificial intelligence (AI) have emerged as two distinct yet profoundly interconnected disciplines for modeling and understanding complex systems. While network science focuses on unraveling the structural and dynamical patterns of interconnected entities constituting systems, AI leverages computational frameworks to infer patterns, predict behaviors, and automate decision-making. These fields are now converging on a few critical junctures, offering unprecedented opportunities to address longstanding challenges in modeling high-dimensional interactions, and emergent phenomena. In this paper, we introduce a few promising research directions pointing out the convergence of these two paradigms and summarize the contributions of a special session organized by the authors at the 33th European Symposium on Artificial Neural Networks, Computational Intelligence, and Machine Learning, showing advances in the integration of network science methodologies and machine learning-based solutions. Throughout this paper we argue that the synthesis of network science's perspective on connecting topology and system functions through robust and principled tools with ML's adaptive learning capabilities forms a transformative paradigm for complexity research.

The convergence of network science and artificial intelligence presents significant disciplinary challenges that must be addressed to fully realize the potential of this integration [3]. From the network science side, network scientists often deal with large, high-dimensional, and dynamic networks, but traditional algorithms may struggle to scale effectively, leading to inefficiencies in analyzing real-time data from complex systems. The need for scalable algorithms while

maintaining accuracy is paramount. Additionally, conventional measurement techniques may not be suitable for modern heterogeneous systems, which require more nuanced approaches to analyze evolving feature-rich networks [16] including heterogeneous node and edge types as well as multiple layers of interactions and arbitrary feature vectors. The integration of AI into network science necessitates effective data representation methods. Many AI techniques require numerical input formats that may not align with the categorical or topological features typical in network data [1]. Converting these features into suitable formats without losing critical information poses a significant challenge, especially if we need to preserve high-order interactions, essential for capturing the intricate relationships within networks [2]. Recently, approaches accounting for multi-node interactions, group dynamics, and higher-dimensional structures, have been proposed. For instance, simplicial complex embeddings extend beyond simple graphs by representing higher-order structures such as higher-dimensional simplices [31]. These embeddings capture the geometric and topological properties of multi-node interactions, allowing for improved tasks, such as higher-order link prediction and network reconstruction. Graph-based methods have also been extended to incorporate high-order interactions. For instance, Higher-Order Graph Neural Networks (HOGNNs) generalize traditional GNNs by aggregating information not only from immediate neighbors but also from multi-hop neighborhoods [42]. Another notable approach involves motif-based embeddings, where recurring patterns or substructures, i.e. network motifs, are explicitly embedded into the representation learning process [39]. These methods leverage the structural significance of motifs to better understand the functional and relational properties of networks. Finally, processes occurring on networks frequently display nonlinear dynamics and emergent behaviors that challenge the accuracy of traditional modeling approaches. These systems, characterized by intricate interactions and feedback loops, often exhibit behaviors that cannot be predicted by simply extrapolating from the properties of individual components. Recent works have emphasized the necessity of incorporating AI's adaptive learning capabilities to address these shortcomings. For example, graph neural networks (GNNs) have demonstrated their effectiveness in modeling emergent phenomena in social and biological networks by learning complex interaction patterns and capturing nonlinear dependencies between network entities, enabling more accurate predictions of phenomena such as information diffusion [4] or protein interactions. Furthermore, research utilizing AI-driven agent-based models is providing insights into the emergence of collective behaviors in social media such as the formation of communities and echo chambers, or the diffusion of possibly harmful content [34].

AI's application to graph-based data confronts several pressing challenges, notably data sparsity, augmentation, interpretability, and validation. Sparse connectivity in many real-world graphs can inhibit model performance and exacerbate difficulties in acquiring sufficiently large or labeled datasets. To address this issue, graph augmentation techniques [47], such as selective edge removal or node addition, have shown promise, yet a comprehensive framework that preserves essential topological features is still lacking [27]. High-order message passing further complicates matters by demanding representations that capture intricate multi-node interactions without losing critical local information. Interpretability and statistical validation also remain key concerns, particularly for temporal graphs [6]. Although post-hoc explanation methods and attention mechanisms have been proposed to elucidate GNN decisions, they often fail to account for evolving network structures. New validation methodologies are similarly needed to assess the reliability of models over time. Here, network science offers valuable insights: tools such as community detection and higherorder interaction analyses [13] can enhance both interpretability and robustness. For instance, pooling strategies that integrate community structures help reduce network complexity while preserving essential information. By systematically evaluating these approaches across diverse temporal graph tasks, researchers can pave the way for more transparent and statistically sound graph-based AI systems.

The scope of the paper encompasses recent advancements in interpretability for temporal graph learning methods, causality-aware graph learning for time series data, representation learning for rich network structures, and graph representation learning aimed at influence maximization. Moreover, the paper introduces the four contributions of the special session with a special focus on interpretable machine-learning approaches, the usage of temporal regularizers to enhance factorization-based temporal knowledge graph completion; hyperbolic representation learning methodologies for node classification, and novel activation functions to create topology-aware neural networks. Collectively, these contributions aim to show advances toward the integration of AI on graphs and network science methodologies.

2 Bridging Network Science and AI

There are multiple areas where a strong interaction between network science and AI is promising: network science supports the design and understanding of AI methods, AI offers new tools to tackle problems in network science, or, vice versa. This section presents key areas of interaction between network science and AI, outlining current literature and promising directions for future research.

2.1 Interpretability of temporal graph learning methods

Temporal Graph Learning (TGL) is a fast-growing field that aims to learn, predict, and understand evolving networks. In recent years, several machinelearning-based solutions have been proposed for forecasting links and node properties on temporal networks, based on common tools such as graph neural networks (GNNs), transformers, or temporal random walks [17, 24, 9]. Temporal Graph Benchmark (TGB)¹ - a widely adopted collection of benchmarks for machine learning on temporal networks - shows that most of these solutions exhibit high variability in the performance over different datasets and tasks.

¹https://tgb.complexdatalab.com/, February 2025

Due to the challenging interpretability of deep learning models, the underlying rationales behind this behavior remain largely opaque. In this context, temporal network theory, network science metrics, and temporal graph mining tools can be leveraged to enhance the interpretability of TGL models. For instance, two works [40, 11] analyze the temporal edge re-occurrence in the TGB datasets, highlighting that deep learning models may vary their performance due to the level of re-occurrence/novelty of the edges in the datasets, depending on their memorization or inductive reasoning capabilities. Dileo et al. [8] tested wellknown network science heuristics for link prediction on TGB. Results show that simple heuristics can reach comparable results with some state-of-the-art techniques and, thanks to their interpretability, give insights into the network being studied, such as the importance of distinguishing link structural roles over time. In the context of explainability techniques for temporal GNNs, Chen et al. [6] proposes a new metric to evaluate the cohesive level of the explanations, based on the temporal proximity of adjacent events, highlighting that most of the models fail to leverage such temporal motifs. Through these examples, it becomes evident that integrating network science principles into AI development can shed light on the inner workings of TGL models. Future research should focus on developing dedicated explainability techniques tailored to temporal graph learning models. Moreover, a key direction is to integrate insights and domain knowledge from network science and temporal graph theory into the design of neural models, for instance, by biasing node and edge representations or integrating temporal motifs and structural constraints to enhance both interpretability and predictive performance. Conversely, advancing the interpretability of temporal graph learning models can also benefit network science research. Compared to traditional shallow embeddings or engineered feature vectors, temporal node embeddings derived from deep learning models may provide richer information on the evolving behavior of nodes, helping to detect distribution shifts and abrupt variations caused by shocking events and mirroring their effect in the learned representation.

2.2 Causality-Aware Deep Learning for Temporal Graph Data

Over the past ten years, the network science community has intensively studied time series data on graphs that can be modeled as *temporal graphs* or *temporal networks* [15]. Apart from aspects like bursty activity patterns, temporal correlations in the activation sequence of edges that influence *time-respecting paths* [18] have been identified as one important aspect that influences the *causal topology* of temporal graphs, i.e. which nodes can influence each other based on the *arrow of time*. In a nutshell, in a temporal graph with two time-stamped edges (a, b; t) and (b, c; t') occurring instantaneously at time stamps t and t', node a can only causally influence node c via a time-respecting path through b iff t < t', i.e. iff the edge (a, b) occurs *before* edge (b, c). Conversely, due to the arrow of time, nodes a and c cannot causally influence each other if the temporal ordering of the two edges is reversed. Several studies have shown that, due to the presence of temporal correlations in the sequential ordering of edges and the

directionality of the arrow of time, real-world temporal graphs exhibit complex causal topologies that considerably change the evolution of dynamical processes like diffusion, epidemic spreading or percolation, community structures, or the centrality of nodes [37, 35, 22, 36]. Despite these findings, this important class of patterns in temporal graph data is often neglected by recent deep learning architectures for temporal graphs, which often aggregate time-stamped edges within batches [24]. These architectures allow us to model patterns in evolving graph representations that are learned in an end-to-end fashion. However, they typically discard information on the microscopic temporal ordering of edges that constitutes the causal topology of temporal graphs and thus influences dynamical processes, node centralities, and community structures. Addressing this gap, recent works at the intersection of deep learning and network science have shown how higher-order graph representations originally developed in network science [20] can be integrated with neural message passing architectures to facilitate prediction tasks in temporal graphs [33, 13]. However, much remains to be explored to further develop causality-aware deep learning methods for temporal graph data.

2.3 Representation learning for rich network structures

When designing AI methods for network science, a major challenge is to develop methods that can properly account for network models that may go beyond classical topologies by integrating more information into the network. Depending on the specific task, feature-rich network models [16] may include heterogeneous node or edge types, multiple layers, and arbitrary attribute vectors associated with nodes and/or edges. In these cases, representation learning methods based on GNN architectures are required to explicitly consider this broader context while producing embedded representations and performing optimization for a specific task at hand. Among the models that extend classic topologies, a widely used one is that of multilayer networks [19], which can model intra-layer and inter-layer relations, different layer characteristics, as well as node features. Zangari et al. [46] propose a GNN framework for representation learning and semi-supervised classification in multilayer networks. The proposed framework is devised to be as flexible as possible with respect to specific instances of multilayer networks since it can take into account an arbitrary number of layers and intra-layer and inter-layer connections between nodes, as well as the presence of side information (i.e., attribute vectors associated to the entities, with possible partial knowledge over the entity set). This represents a major improvement to other approaches dealing with multi-relational networks, such as GrAMME [38], that can only take into account *multiplex* topologies, e.g., where arbitrary interlayer connections are not allowed, but only coupling edges among different instances of a same entity. However, real-world systems may not only include several heterogeneous interacting entities but also be inherently dynamic, with relations and entity sets that continuously evolve. Intending to propose a tool capable of dealing with such scenarios, Martirano et al. [26] proposed DyHANE (Dynamic Heterogeneous Attributed Network Embedding), a framework that

exploits the Graph Continual Learning (GCL) paradigm for a data-incremental multi-class node classification task on Heterogeneous Information Networks (HINs). DyHANE identifies a representative sample of multi-typed nodes as training set and updates the parameters of a GNN module, to generate up-to-date representations for all nodes in the network. The main idea is to avoid completely re-training the model when only minor changes affect the network, and to reduce the risk of catastrophic forgetting [48], by extending strategies like *experience replay* and *generative replay* [44, 45] to the HIN case. These strategies are mainly based on the ability to store a representative subset of past data or of synthetic prototypes in a memory buffer, to consolidate acquired knowledge about existing patterns during the dynamic evolution of the network.

2.4 Graph representation learning for influence maximization

Identifying influential spreaders—individuals capable of effectively disseminating information—is essential for deepening our understanding of spreading dynamics in social systems [25]. The problem lies at the heart of numerous real-world applications, including computational epidemiology, viral marketing, media polarization, and misinformation detection. In a particular instance of the problem, known as *influence maximization*, the goal is to find the optimal subset of k nodes that maximize the influence spread under a stochastic information diffusion model that simulates how influence propagates over the graph (e.g., the independent cascade model). The problem has been extensively studied from a combinatorial optimization perspective, with a primary emphasis on addressing time complexity challenges. With the recent advances in graph machine learning, the research community started investigating how representation learning techniques can be leveraged to address the influence learning and maximization problems. A particular challenge here is how real diffusion cascades (e.g., retweets in \mathbb{X}) can be utilized in the process, avoiding overfitting due to the large number of parameters (i.e., proportional to the number of edges) [12] as well as assumptions imposed by the diffusion model [4]. In this direction, Panagopoulos et al. [29] introduced IMINFECTOR, an influence maximization framework devoid of diffusion models that relies on representations learned from real information cascades. The model consists of a multi-task neural network that simultaneously learns embeddings for nodes initiating cascades (influencer vectors) and those participating in them (susceptible vectors) while also predicting the cascade size. These embedding vectors encode diffusion probabilities among nodes, allowing a greedy approximation algorithm to be employed with theoretical guarantees. To further improve expressiveness, Ling et al. [23] introduced DeepIM, a model that relies on GNNs to learn representations of seed sets and, subsequently, to define a seed selection policy to maximize the influence spread. GLIE [30] employs a similar methodology, resulting in an inductive model trained on small graphs capable of generalizing on significantly larger test graphs. More recently, these approaches have been extended to learn the underlying diffusion model through cellular sheaf GNNs, including the DeepSN model [14].

Let us mention here that this line of work on influence maximization is

part of a broader effort to leverage graph learning for addressing combinatorial optimization problems [5]. Similar approaches have been applied to other well-known network science tasks, such as graph clustering and community detection, to optimize objectives like modularity [43] or graph cuts [10]. This further demonstrates the potential of AI-driven solutions, aiming to balance efficiency, generalization, and graph structure expressiveness.

3 Special Session's Contributions

This year's special session comprises original contributions ranging on a set of diverse topics on the intersection between network science and artificial intelligence:

- Lange-Geisler *et al.* [21] propose an interpretable machine-learning approach to assess the reliability of network-structured systems. Their central idea is to use a classification scheme to identify prototype vectors in feature space representing the (un-)reliability levels. For this, they apply the Generalized Matrix Relevance Learning Vector Quantization (GMLVQ).
- Dileo *et al.* [7] systematically analyzeseveral choices of temporal regularizers to enhance the performance of factorization-based temporal knowledge graph completion methods. Their work shows that by carefully tuning simple tensor factorization models, they can reach comparable performance with other competitors, enabling several applications useful in network science and graph mining, such as producing more interpretable results or scaling to very large graphs, without substantial sacrifices in performance.
- Pogány *et al.* [32] propose a methodology based on hyperbolic representation learning for node classification in protein-protein interaction (PPI) networks. Each PPI network is associated with a specific tissue, which is organized in a hierarchical tree-like structure. The method leverages this structure to allow multiple representations of the same protein in adjacent layers of the hierarchy to share similar features. Specifically, the author extends OhmNet [49] to work in a hyperbolic space, leveraging the Lorentz model. Results show, on average, an increase in performance compared to the Euclidean counterpart.
- Snopov *et al.* [41] propose novel activation functions that enhance the ability of neural networks to manipulate data topology during training. Their work investigates whether non-standard activation functions can be used as the basis of more topology-aware neural networks. The work builds on previous research demonstrating that neural networks elaborate data for classification by progressive simplification of the topology of the data [28].

4 Conclusion

In conclusion, this paper highlights the synergistic relationship between network science and AI, showcasing how their integration can address complex challenges in data analysis and modeling. By outlining promising research directions, and summarizing contributions presented at the ESANN 2025 special session, we aim to highlight the potential of combining network science and AI. The findings underscore the need for further exploration of these hybrid approaches to enhance interpretability, validation, and overall performance across various application domains.

References

- A. Baptista, R. J. Sánchez-García, A. Baudot, and G. Bianconi. Zoo guide to network embedding. Journal of Physics: Complexity, 4(4):042001, 2023.
- [2] G. Bianconi. Higher-order networks. Cambridge University Press, 2021.
- [3] C. Blöcker, M. Rosvall, I. Scholtes, and J. D. West. Insights from network science can advance deep graph learning, 2025.
- [4] S. Bourigault, S. Lamprier, and P. Gallinari. Representation learning for information diffusion through social networks: an embedded cascade model. In <u>WSDM</u>, 2016.
- [5] Q. Cappart, D. Chételat, E. B. Khalil, A. Lodi, C. Morris, and P. Velčcković. Combinatorial optimization and reasoning with graph neural networks. <u>Journal of</u> Machine Learning Research, 24(130):1–61, 2023.
- [6] J. Chen and R. Ying. Tempme: Towards the explainability of temporal graph neural networks via motif discovery. In NeurIPS, 2023.
- [7] M. Dileo, P. Minervini, M. Zignani, and S. Gaito. Enhancing neural link predictors for temporal knowledge graphs with temporal regularisers. In ESANN, 2025.
- [8] M. Dileo and M. Zignani. Link prediction heuristics for temporal graph benchmark. In ESANN, 2024.
- M. Dileo, M. Zignani, and S. Gaito. Temporal graph learning for dynamic link prediction with text in online social networks. <u>Machine Learning</u>, 113(4):2207–2226, Apr 2024.
- [10] A. Duval and F. Malliaros. Higher-order clustering and pooling for graph neural networks. In CIKM, 2022.
- [11] J. Gastinger, C. Meilicke, F. Errica, T. Sztyler, A. Schuelke, and H. Stuckenschmidt. History repeats itself: A baseline for temporal knowledge graph forecasting. 2024.
- [12] A. Goyal, F. Bonchi, and L. V. S. Lakshmanan. A data-based approach to social influence maximization. Proc. VLDB Endow., 5(1):73–84, 2011.
- [13] F. Heeg and I. Scholtes. Using time-aware graph neural networks to predict temporal centralities in dynamic graphs. In <u>Proceedings of the 38th Conference</u> on Advances in Neural Information Processing Systems (NeurIPS 2024), 2024.
- [14] A. Hevapathige, Q. Wang, and A. N. Zehmakan. DeepSN: A sheaf neural framework for influence maximization. In AAAI, 2025.
- [15] P. Holme and J. Saramäki. Temporal networks. <u>Physics reports</u>, 519(3):97–125, 2012.
- [16] R. Interdonato, M. Atzmueller, S. Gaito, R. Kanawati, C. Largeron, and A. Sala. Feature-rich networks: going beyond complex network topologies. <u>Appl. Netw.</u> Sci., 4(1):4:1–4:13, 2019.

ESANN 2025 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 23-25 April 2025, i6doc.com publ., ISBN 9782875870933. Available from http://www.i6doc.com/en/.

- [17] S. M. Kazemi, R. Goel, K. Jain, I. Kobyzev, A. Sethi, P. Forsyth, and P. Poupart. Representation learning for dynamic graphs: A survey. J. Mach. Learn. Res., 21:70:1–70:73, 2020.
- [18] D. Kempe, J. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. In Proceedings of the thirty-second annual ACM symposium on Theory of computing, pages 504–513, 2000.
- [19] M. Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter. Multilayer networks. J. Complex Networks, 2(3):203–271, 2014.
- [20] R. Lambiotte, M. Rosvall, and I. Scholtes. From networks to optimal higher-order models of complex systems. Nature physics, 15(4):313–320, 2019.
- [21] M. Lange-Geisler, K. Dohmen, and T. Villmann. Learning of probability estimates for system and network reliability analysis by means of matrix learning vector quantization. In ESANN, 2025.
- [22] H. H. Lentz, T. Selhorst, and I. M. Sokolov. Unfolding accessibility provides a macroscopic approach to temporal networks. <u>Physical review letters</u>, 110(11):118701, 2013.
- [23] C. Ling, J. Jiang, J. Wang, M. T. Thai, L. Xue, J. Song, M. Qiu, and L. Zhao. Deep graph representation learning and optimization for influence maximization. In ICML, 2023.
- [24] A. Longa, V. Lachi, G. Santin, M. Bianchini, B. Lepri, P. Lio, F. Scarselli, and A. Passerini. Graph neural networks for temporal graphs: State of the art, open challenges, and opportunities, 2023.
- [25] F. D. Malliaros, M.-E. G. Rossi, and M. Vazirgiannis. Locating influential nodes in complex networks. Scientific Reports, 6:19307, 2016.
- [26] L. Martirano, D. Ienco, R. Interdonato, and A. Tagarelli. Dyhane: dynamic heterogeneous attributed network embedding through experience node replay. Appl. Netw. Sci., 9(1):30, 2024.
- [27] C. Morris, F. Frasca, N. Dym, H. Maron, I. I. Ceylan, R. Levie, D. Lim, M. M. Bronstein, M. Grohe, and S. Jegelka. Position: Future directions in the theory of graph machine learning. In R. Salakhutdinov, Z. Kolter, K. Heller, A. Weller, N. Oliver, J. Scarlett, and F. Berkenkamp, editors, <u>Proceedings of the 41st International Conference on Machine Learning</u>, volume 235 of <u>Proceedings of Machine Learning Research</u>, pages 36294–36307. PMLR, 21–27 Jul 2024.
- [28] G. Naitzat, A. Zhitnikov, and L.-H. Lim. Topology of deep neural networks. <u>J.</u> Mach. Learn. Res., 21(1), Jan. 2020.
- [29] G. Panagopoulos, F. D. Malliaros, and M. Vazirgiannis. Multi-task learning for influence estimation and maximization. <u>IEEE Trans. Knowl. Data Eng.</u>, 34(9):4398– 4409, 2022.
- [30] G. Panagopoulos, N. Tziortziotis, M. Vazirgiannis, and F. D. Malliaros. Maximizing influence with graph neural networks. In ASONAM, 2023.
- [31] S. Piaggesi, A. Panisson, and G. Petri. Effective higher-order link prediction and reconstruction from simplicial complex embeddings. In <u>Learning on Graphs</u> Conference, pages 55–1. PMLR, 2022.
- [32] D. Pogány and P. Antal. Hyperbolic representation learning in multi-layer tissue networks. In ESANN, 2025.

ESANN 2025 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 23-25 April 2025, i6doc.com publ., ISBN 9782875870933. Available from http://www.i6doc.com/en/.

- [33] L. Qarkaxhija, V. Perri, and I. Scholtes. De bruijn goes neural: Causality-aware graph neural networks for time series data on dynamic graphs. In <u>Learning on</u> Graphs Conference, pages 51–1. PMLR, 2022.
- [34] G. Rossetti, M. Stella, R. Cazabet, K. Abramski, E. Cau, S. Citraro, A. Failla, R. Improta, V. Morini, and V. Pansanella. Y social: an llm-powered social media digital twin. arXiv preprint arXiv:2408.00818, 2024.
- [35] M. Rosvall, A. V. Esquivel, A. Lancichinetti, J. D. West, and R. Lambiotte. Memory in network flows and its effects on spreading dynamics and community detection. Nature communications, 5(1):4630, 2014.
- [36] I. Scholtes. When is a network a network? multi-order graphical model selection in pathways and temporal networks. In <u>Proceedings of the 23rd ACM SIGKDD</u> <u>international conference on knowledge discovery and data mining</u>, pages 1037– 1046, 2017.
- [37] I. Scholtes, N. Wider, R. Pfitzner, A. Garas, C. J. Tessone, and F. Schweitzer. Causality-driven slow-down and speed-up of diffusion in non-markovian temporal networks. Nature communications, 5(1):5024, 2014.
- [38] U. S. Shanthamallu, J. J. Thiagarajan, H. Song, and A. Spanias. GrAMME: Semisupervised Learning Using Multilayered Graph Attention Models. <u>IEEE</u> Transactions on Neural Networks and Learning Systems, 31(10):3977–3988, 2020.
- [39] P. Shao, Y. Yang, S. Xu, and C. Wang. Network embedding via motifs. <u>ACM</u> Transactions on Knowledge Discovery from Data (TKDD), 16:1 – 20, 2021.
- [40] R. Shirzadkhani, S. Huang, E. Kooshafar, R. Rabbany, and F. Poursafaei. Temporal graph analysis with TGX. In WSDM, pages 1086–1089. ACM, 2024.
- [41] P. Snopov and O. Musin. Topology-aware activation functions in neural networks. In ESANN, 2025.
- [42] H. Tian and R. Zafarani. Higher-order networks representation and learning: A survey. SIGKDD Explor. Newsl., 26(1):1–18, July 2024.
- [43] A. Tsitsulin, J. Palowitch, B. Perozzi, and E. Müller. Graph clustering with graph neural networks. Journal of Machine Learning Research, 24(127):1–21, 2023.
- [44] J. Wang, G. Song, Y. Wu, and L. Wang. Streaming graph neural networks via continual learning. <u>CoRR</u>, abs/2009.10951, 2020.
- [45] J. Wang, W. Zhu, G. Song, and L. Wang. Streaming graph neural networks with generative replay. In <u>Proceedings of the 28th ACM SIGKDD Conference on</u> Knowledge Discovery and Data Mining, pages 1878–1888, 2022.
- [46] L. Zangari, R. Interdonato, A. Caliò, and A. Tagarelli. Graph convolutional and attention models for entity classification in multilayer networks. <u>Appl. Netw. Sci.</u>, 6(1):87, 2021.
- [47] T. Zhao, W. Jin, Y. Liu, Y. Wang, G. Liu, S. Günnemann, N. Shah, and M. Jiang. Graph data augmentation for graph machine learning: A survey. <u>IEEE Data Eng.</u> Bull., 2023.
- [48] F. Zhou and C. Cao. Overcoming catastrophic forgetting in graph neural networks with experience replay. In <u>Proceedings of the AAAI Conference on Artificial</u> Intelligence, volume 35, pages 4714–4722, 2021.
- [49] M. Zitnik and J. Leskovec. Predicting multicellular function through multi-layer tissue networks. Bioinformatics, 33(14):190–198, 2017.