

Opportunities for graph learning in robotics

Original

Opportunities for graph learning in robotics / Pistilli, Francesca; Averta, Giuseppe. - (2023). (Intervento presentato al convegno I-RIM 3D 2023).

Availability:

This version is available at: 11583/2982854 since: 2023-10-09T09:59:37Z

Publisher:

I-RIM

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Opportunities for graph learning in robotics

1st Francesca Pistilli

DAUIN, Polytechnic of Turin, Turin, Italy
francesca.pistilli@polito.it

2nd Giuseppe Averta

DAUIN, Polytechnic of Turin, Turin, Italy
giuseppe.averta@polito.it

Abstract—In the last few years, robotics highly benefited from the use of machine and deep learning to process data stream captured by robots during their tasks. Yet, encoding data in grids (images) or vectors (time-series) significantly limits the type of data that can be processed to euclidean only. To unlock the potential of deep learning also to unstructured data, such as point clouds or functional relations, a rising - yet under-explored - approach lies on the use of graph neural networks (GNNs). With this manuscript, we intend to deliver a brief introduction to GNNs for robotics applications, together with a concise revision of notable applications in the field, with the aim of fostering the use of this learning strategy in a wider context and highlighting potential future research directions.

Index Terms—Graph Neural Networks, Robotics, Deep Learning, Human-Machine Interaction

I. INTRODUCTION

Robotics represent the perfect ultimate application for deep learning methods, where theoretical methods can meet realistic and challenging applications. Recently, in robotic learning has emerged the necessity to use a richer knowledge representation, able to represent three-dimensional data, functional and geometrical representation in contrast to a simple embedding of bi-dimensional or temporal sequence representations. A promising approach, inspired from the signal processing and computer vision communities, rely of graph representations. Graph neural networks are powerful mathematical tools for knowledge representation, able to encode structured and unstructured information through nodes which model entities, and edges that encode their spatial, temporal or functional relationship. In computer vision, learning on graphs has emerged as an effective approach to deal with data lying on irregular domains, such as point clouds [1], [2], to uncover non-local similarities in the data [3], and video understanding [4], [5]. In the last few years, these methods have gained the attention of the robotic community, opening interesting perspectives for novel tasks and applications, ranging from grasping and manipulation, where three-dimensional data are predominant, to robot-object interaction, where graphs can be used to model the environment, the objects and their relationships. Notably, GNNs showed promising results also in the context of reinforcement learning, showing interesting abilities in predicting future states based on current state and action signals. One of the main limitation of deep reinforcement learning is the limited generalisation ability, and GNNs are able to naturally overcome this issue and learn robust policies across scenarios [6]–[8].

II. PRELIMINARIES ON GRAPH NEURAL NETWORKS

Learning on graph structures has gained significant relevance in the community, thanks to their ability to process data lying on irregular domains, spatio-temporal and functional relationships [9]–[11]. A graph is defined as a set of nodes, the entities of the space under observation, and edges, the connections between nodes. Each node and edge can be associated with a feature vector to encode information about elements and their relationship. GNNs process data represented on graph structures by performing a local share of information, also known as message-passing, in order to update the elements features. A graph convolution operation can be formulated in two different domains, spectral or spatial. The first family of methods [12]–[16] usually exploits graph Fourier transform, eventually complemented with polynomial approximations to reduce the computational burden [14], [15]. Among these, it is worth mentioning the Graph Convolutional Network (GCN) [15], which demonstrated notable results for semi-supervised problems. However, a critical limitation of this formulation is the inability to generalise the learned filters, computed over the spectrum of the graph Laplacian, to a variable graph structure. The second class of approaches defines the graph-convolution operation in the spatial domain. In this case, the graph convolution is defined as a local, i.e. computed over a neighbourhood, weighted aggregation of signals. Since the interaction between nodes is local, this formulation is suitable for any type of signal that can be defined over a graph, even with a variable graph structure. Several definitions have been presented in the literature, among which it is worth mentioning [1], [2], [18]–[24]. Typically, the formulations differ on the computation of the weights used in the aggregation. For a compact overview, the reader is referred to Fig. 1 where we show different graph convolution formulations.

III. ROBOT APPLICATIONS

We identified two main macro-problems in which GNNs represent a valid learning tools. The first, which we refer to as "Single-Agent Systems", collects cases where the graph is used to model only one intelligent system (single-robot or human), or passive body (objects modelling). One of the most important tasks in this scenario is modelling of bodies. When we interact with an object very likely it will suffer a non-linear deformation, and modelling its dynamic and deformation poses several challenges such as the high-dimensionality of the configuration space, complex dynamic of materials and self-occlusions. One promising approach [25], [26] consists in

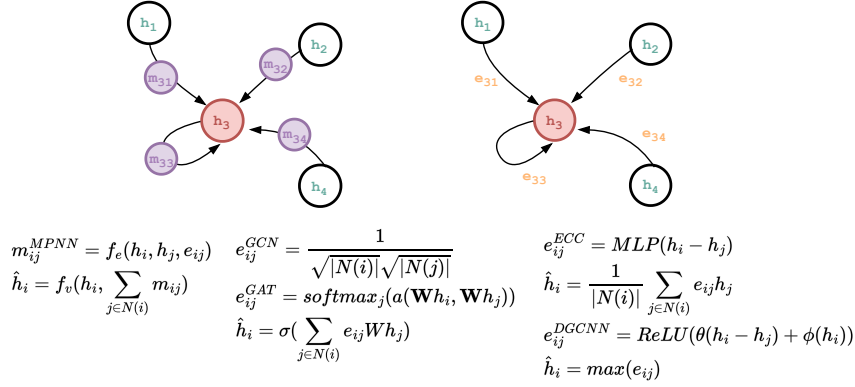


Fig. 1. Comparison of different information sharing mechanisms. MPNN [17] implements message passing conditioned by edge attributes. GCN [15] implements a simple aggregation of the k-nn nodes. GAT [18] exploits self-attention to weight the contributions of the k-nn. ECC [1] implements an edge-dependent weighting function to weight the contribution of each node in the neighbourhood. DGCNN [2] aggregates by means of learned edge features.

describing the body as a collection of keypoints. This approximated representation is suitable to a graph construction, where each keypoint is a node and the edges represent the spatial and kinetic relationships between them, able to describe the dynamic behaviour of the body. Recently, graph-based learning control policies from deformable objects keypoints has gained attention [26]–[31] thanks to their effective low-dimensional representation. One challenge of modelling deformable objects rely on the object connectivity, due to the large number of degree of freedom and self-occlusion. Graphs can easily overcome this issue by directly learning the connectivity from the visual data [32], [33]. Furthermore, graphs represent one of the most suitable choice to process three-dimensional visual data, such as point cloud and meshes, and recently this type of data - especially mesh - have gained attention for particles and keypoints representation, promoting even more the use of GNNs in the context of body modelling [32]–[35]. Similarly to keypoints object representations, graphs can be used to model the robot itself, exploiting the natural discrete graph structure of human, humanoid or animal bodies, where nodes are joints and edges their physical dependencies [36], [37]. Once we have learned how to model robots and objects it is important to learn how to interact safely and effectively with them. Albeit being a research problem for decades [38], [39], in robotic grasping and manipulation several problems are still open, such as learning to interact with objects in motion or in clutter [40]–[42] with partial observability, novel objects and human-robot or multi-robot co-manipulation [43]–[46]. In the last few years, learning for graphs has provided interesting results in tackling many open problems in this field. Graphs can efficiently model the environment and the interaction between objects [47], [48], learn semantic global information to build knowledge graphs [49]–[51] or process unstructured data input [52]–[54]. Graph-based structures have been also profitably used to encode spatial relations in complex scenes, where multiple objects are present and/or the environment is only partially observable [48], [55]–[62].

The second macro-problem we identified refers to the modeling of fleets of intelligent systems (a.k.a. Multi-agent systems). This scenario comes with the additional challenge

that an effective coordination and communication between robots is of paramount importance to enable a fruitful cooperation, which is clearly critical for several downstream tasks. Traditionally, this problem has been addressed by exploiting centralised approaches, where all the intelligence and computation complexity is centralised on a common controller node. Such approach has the strong benefits of reasoning on the status of whole set of agents and requiring only little computation on the device. However, it also comes with the limitations of i) a low fault tolerance, and ii) a poor scalability w.r.t. the number of agents [63]. To address these issues, decentralised approaches, where each agent learns its set of actions according to environmental data and shared information from other agents, have gained attention. In this configuration, each robot has to deal with limited observability and partial communication, but GNNs offer generalized and flexible structural representations of the elements to face such limitations. Many notable works [64]–[74] focus the attention on the communication between robots and exploit graph encoding to efficiently model inter-agents relationships. Other methods, instead, propose to encode the information related to the environment topology, in the form of waypoints or point-of-interest [6], [7], [75]–[78]. Works as [6], [7], [75] focus on the spatio-temporal relationship between entities in the environment to build a graph, and use GNNs to learn environmental features for the robot. Differently, [76], [77] propose a more comprehensive topological representation where both robots and environments are directly encoded in the graph.

IV. FUTURE PROSPECTIVES

After a literature analysis, it is evident how graph-based representations opened interesting research perspectives, such as for the modelling of functional, spatial, or temporal relationships between passive or active elements. Interestingly, we observed an exponential increase in the use of this learning method for robotics applications, which at the moment is consolidated as one of the most prominent for complex tasks. For the field to progress to a higher maturity phase, however, we strongly believe that some aspects still deserve additional work. First, the analysis we performed strongly motivates a

more extensive and informed use of graphs as a convenient tool. Yet, it is often used in a naive fashion. Furthermore, graph learning could unlock many novel tasks, going from action recognition to modelling of soft-bodies. We believe that robotics should not only be an harbour for graph learning applications, but also actively contribute in the development of novel techniques that better face real-world problems.

ACKNOWLEDGMENT

This study was carried out within the FAIR - Future Artificial Intelligence Research and received funding from the European Union Next-GenerationEU (PIANO NAZIONALE DI RIPRESA E RESILIENZA (PNRR) – MISSIONE 4 COMPONENTE 2, INVESTIMENTO 1.3 – D.D. 1555 11/10/2022, PE00000013). This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

REFERENCES

- [1] M. Simonovsky and N. Komodakis, "Dynamic edge-conditioned filters in convolutional neural networks on graphs," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 3693–3702.
- [2] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *Acm Transactions On Graphics (tog)*, vol. 38, no. 5, pp. 1–12, 2019.
- [3] D. Valsesia, G. Fracastoro, and E. Magli, "Deep graph-convolutional image denoising," *IEEE Transactions on Image Processing*, vol. 29, pp. 8226–8237, 2020.
- [4] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, and A. Gupta, "Hollywood in homes: Crowdsourcing data collection for activity understanding," in *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*. Springer, 2016, pp. 510–526.
- [5] I. Nagarajan, Y. Li, C. Feichtenhofer, and K. Grauman, "Ego-topo: Environment affordances from egocentric video," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 163–172.
- [6] F. Chen, J. D. Martin, Y. Huang, J. Wang, and B. Englot, "Autonomous exploration under uncertainty via deep reinforcement learning on graphs," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 6140–6147.
- [7] F. Chen, P. Szenher, Y. Huang, J. Wang, T. Shan, S. Bai, and B. Englot, "Zero-shot reinforcement learning on graphs for autonomous exploration under uncertainty," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 5193–5199.
- [8] T. Luo, B. Subagdja, D. Wang, and A.-H. Tan, "Multi-agent collaborative exploration through graph-based deep reinforcement learning," in *2019 IEEE International Conference on Agents (ICA)*. IEEE, 2019, pp. 2–7.
- [9] D. Bacciu, F. Errica, A. Micheli, and M. Podda, "A gentle introduction to deep learning for graphs," *Neural Networks*, vol. 129, pp. 203–221, 2020.
- [10] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip, "A comprehensive survey on graph neural networks," *IEEE transactions on neural networks and learning systems*, vol. 32, no. 1, pp. 4–24, 2020.
- [11] H. Yuan, H. Yu, S. Gui, and S. Ji, "Explainability in graph neural networks: A taxonomic survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [12] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, "Spectral networks and locally connected networks on graphs," *arXiv preprint arXiv:1312.6203*, 2013.
- [13] M. Henaff, J. Bruna, and Y. LeCun, "Deep convolutional networks on graph-structured data," *arXiv preprint arXiv:1506.05163*, 2015.
- [14] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Advances in neural information processing systems*, 2016, pp. 3844–3852.
- [15] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.
- [16] F. Monti, K. Otness, and M. M. Bronstein, "Motifnet: a motif-based graph convolutional network for directed graphs," in *2018 IEEE Data Science Workshop (DSW)*. IEEE, 2018, pp. 225–228.
- [17] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *International conference on machine learning*. PMLR, 2017, pp. 1263–1272.
- [18] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," *6th International Conference on Learning Representations*, 2017.
- [19] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, and M. M. Bronstein, "Geometric deep learning on graphs and manifolds using mixture model cnns," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5115–5124.
- [20] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Advances in Neural Information Processing Systems*, 2017, pp. 1024–1034.
- [21] N. Verma, E. Boyer, and J. Verbeek, "Featnet: Feature-steered graph convolutions for 3d shape analysis," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2598–2606.
- [22] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?" in *International Conference on Learning Representations*, 2018.
- [23] L. Wang, Y. Huang, Y. Hou, S. Zhang, and J. Shan, "Graph attention convolution for point cloud semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10296–10305.
- [24] G. Corso, L. Cavalleri, D. Beaini, P. Liò, and P. Veličković, "Principal neighbourhood aggregation for graph nets," *arXiv preprint arXiv:2004.05718*, 2020.
- [25] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia, "Learning to simulate complex physics with graph networks," in *International conference on machine learning*. PMLR, 2020, pp. 8459–8468.
- [26] Y. Li, J. Wu, R. Tedrake, J. B. Tenenbaum, and A. Torralba, "Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids," *arXiv preprint arXiv:1810.01566*, 2018.
- [27] X. Ma, D. Hsu, and W. S. Lee, "Learning latent graph dynamics for visual manipulation of deformable objects," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 8266–8273.
- [28] C. Wang, Y. Zhang, X. Zhang, Z. Wu, X. Zhu, S. Jin, T. Tang, and M. Tomizuka, "Offline-online learning of deformation model for cable manipulation with graph neural networks," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5544–5551, 2022.
- [29] H. Shi, H. Xu, Z. Huang, Y. Li, and J. Wu, "Robocraft: Learning to see, simulate, and shape elasto-plastic objects with graph networks," *arXiv preprint arXiv:2205.02909*, 2022.
- [30] Y. Deng, C. Xia, X. Wang, and L. Chen, "Graph-transporter: A graph-based learning method for goal-conditioned deformable object rearranging task," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2022, pp. 1910–1916.
- [31] Y. Deng, X. Wang, and L. Chen, "Learning visual-based deformable object rearrangement with local graph neural networks," *Complex & Intelligent Systems*, pp. 1–14, 2023.
- [32] X. Lin, Y. Wang, Z. Huang, and D. Held, "Learning visible connectivity dynamics for cloth smoothing," in *Conference on Robot Learning*. PMLR, 2022, pp. 256–266.
- [33] K. Mo, Y. Deng, C. Xia, and X. Wang, "Learning language-conditioned deformable object manipulation with graph dynamics," *arXiv preprint arXiv:2303.01310*, 2023.
- [34] T. Pfaff, M. Fortunato, A. Sanchez-Gonzalez, and P. W. Battaglia, "Learning mesh-based simulation with graph networks," *arXiv preprint arXiv:2010.03409*, 2020.
- [35] I. Huang, Y. Narang, R. Bajcsy, F. Ramos, T. Hermans, and D. Fox, "Defgraspnet: Grasp planning on 3d fields with graph neural nets," *arXiv preprint arXiv:2303.16138*, 2023.
- [36] T. Wang, R. Liao, J. Ba, and S. Fidler, "Nervenet: Learning structured policy with graph neural networks," in *Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada*, vol. 30, 2018.
- [37] M. Oliva, S. Banik, J. Josifovski, and A. Knoll, "Graph neural networks for relational inductive bias in vision-based deep reinforcement learning of robot control," in *2022 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2022, pp. 1–9.

- [38] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *Proceedings 2000 ICRA. Millennium conference. IEEE international conference on robotics and automation. Symposia proceedings (Cat. No. 00CH37065)*, vol. 1. IEEE, 2000, pp. 348–353.
- [39] K. Kleeberger, R. Bormann, W. Kraus, and M. F. Huber, "A survey on learning-based robotic grasping," *Current Robotics Reports*, vol. 1, pp. 239–249, 2020.
- [40] B. Wen, W. Lian, K. Bekris, and S. Schaal, "Catgrasp: Learning category-level task-relevant grasping in clutter from simulation," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6401–6408.
- [41] L. Berscheid, P. Meißner, and T. Krüger, "Robot learning of shifting objects for grasping in cluttered environments," in *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2019, pp. 612–618.
- [42] M. R. Dogar, K. Hsiao, M. Ciocarlie, and S. Srinivasa, *Physics-based grasp planning through clutter*. MIT Press: Cambridge, MA, USA, 2012.
- [43] X. Li, G. Chi, S. Vidas, and C. C. Cheah, "Human-guided robotic co-manipulation: Two illustrative scenarios," *IEEE Transactions on Control Systems Technology*, vol. 24, no. 5, pp. 1751–1763, 2016.
- [44] L. Pernel, N. Tsagarakis, D. Caldwell, and A. Ajoudani, "Adaptation of robot physical behaviour to human fatigue in human-robot co-manipulation," in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2016, pp. 489–494.
- [45] F. Dimeas and N. Aspragathos, "Reinforcement learning of variable admittance control for human-robot co-manipulation," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 1011–1016.
- [46] M. L. Elwin, B. Strong, R. A. Freeman, and K. M. Lynch, "Human-multirobot collaborative mobile manipulation: the omnid mobocobots," *IEEE Robotics and Automation Letters*, vol. 8, no. 1, pp. 376–383, 2022.
- [47] F. Xie, A. Chowdhury, M. De Paolis Kaluza, L. Zhao, L. Wong, and R. Yu, "Deep imitation learning for bimanual robotic manipulation," *Advances in neural information processing systems*, vol. 33, pp. 2327–2337, 2020.
- [48] Y. Lin, A. S. Wang, E. Undersander, and A. Rai, "Efficient and interpretable robot manipulation with graph neural networks," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2740–2747, 2022.
- [49] A. Murali, W. Liu, K. Marino, S. Chernova, and A. Gupta, "Same object, different grasps: Data and semantic knowledge for task-oriented grasping," in *Conference on Robot Learning*, 2020.
- [50] J. H. Kwak, J. Lee, J. J. Whang, and S. Jo, "Semantic grasping via a knowledge graph of robotic manipulation: A graph representation learning approach," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9397–9404, 2022.
- [51] L. Collodi, D. Bacciu, M. Bianchi, and G. Avverta, "Learning with few examples the semantic description of novel human-inspired grasp strategies from rgb data," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2573–2580, 2022.
- [52] M. Neumann, P. Moreno, L. Antanas, R. Garnett, and K. Kersting, "Graph kernels for object category prediction in task-dependent robot grasping," in *Online proceedings of the eleventh workshop on mining and learning with graphs*, 2013, pp. 0–6.
- [53] A. Alliegro, M. Rudorfer, F. Frattin, A. Leonardis, and T. Tommasi, "End-to-end learning to grasp via sampling from object point clouds," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9865–9872, 2022.
- [54] A. Iriondo, E. Lazkano, and A. Ansuategi, "Affordance-based grasping point detection using graph convolutional networks for industrial bin-picking applications," *Sensors*, vol. 21, no. 3, p. 816, 2021.
- [55] G. Zuo, J. Tong, Z. Wang, and D. Gong, "A graph-based deep reinforcement learning approach to grasping fully occluded objects," *Cognitive Computation*, vol. 15, no. 1, pp. 36–49, 2023.
- [56] X. Lou, Y. Yang, and C. Choi, "Learning object relations with graph neural networks for target-driven grasping in dense clutter," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 742–748.
- [57] Y. Huang, A. Conkey, and T. Hermans, "Planning for multi-object manipulation with graph neural network relational classifiers," *arXiv preprint arXiv:2209.11943*, 2022.
- [58] M. Ding, Y. Liu, C. Yang, and X. Lan, "Visual manipulation relationship detection based on gated graph neural network for robotic grasping," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1404–1410.
- [59] R. Li, A. Jabri, T. Darrell, and P. Agrawal, "Towards practical multi-object manipulation using relational reinforcement learning," in *2020 IEEE international conference on robotics and automation (icra)*. IEEE, 2020, pp. 4051–4058.
- [60] I. Kapelyukh and E. Johns, "My house, my rules: Learning tidying preferences with graph neural networks," in *Conference on Robot Learning*. PMLR, 2022, pp. 740–749.
- [61] M. Wilson and T. Hermans, "Learning to manipulate object collections using grounded state representations," in *Conference on Robot Learning*. PMLR, 2020, pp. 490–502.
- [62] T. Silver, R. Chitnis, A. Curtis, J. B. Tenenbaum, T. Lozano-Pérez, and L. P. Kaelbling, "Planning with learned object importance in large problem instances using graph neural networks," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 13, 2021, pp. 11 962–11 971.
- [63] A. Khan, A. Ribeiro, V. R. Kumar, and A. G. Francis, "Graph neural networks for motion planning," *ArXiv*, vol. abs/2006.06248, 2020.
- [64] C. Sun, M. Shen, and J. P. How, "Scaling up multiagent reinforcement learning for robotic systems: Learn an adaptive sparse communication graph," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11 755–11 762.
- [65] E. Tolstaya, F. Gama, J. Paulos, G. Pappas, V. Kumar, and A. Ribeiro, "Learning decentralized controllers for robot swarms with graph neural networks," in *Conference on robot learning*. PMLR, 2020, pp. 671–682.
- [66] Z. Wang and M. Gombolay, "Learning scheduling policies for multi-robot coordination with graph attention networks," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4509–4516, 2020.
- [67] J. Blumenkamp and A. Prorok, "The emergence of adversarial communication in multi-agent reinforcement learning," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 1394–1414.
- [68] R. Kortvelesy and A. Prorok, "Modgnn: Expert policy approximation in multi-agent systems with a modular graph neural network architecture," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 9161–9167.
- [69] L. Clark, J. Galante, B. Krishnamachari, and K. Psounis, "A queue-stabilizing framework for networked multi-robot exploration," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2091–2098, 2021.
- [70] Z. Wang, C. Liu, and M. Gombolay, "Heterogeneous graph attention networks for scalable multi-robot scheduling with temporospatial constraints," *Autonomous Robots*, pp. 1–20, 2022.
- [71] E. Sebastian, T. Duong, N. Atanasov, E. Montijano, and C. Sagues, "Lemurs: Learning distributed multi-robot interactions," *arXiv preprint arXiv:2209.09702*, 2022.
- [72] F. Gama, Q. Li, E. Tolstaya, A. Prorok, and A. Ribeiro, "Synthesizing decentralized controllers with graph neural networks and imitation learning," *IEEE Transactions on Signal Processing*, vol. 70, pp. 1932–1946, 2022.
- [73] W. Gosrich, S. Mayya, R. Li, J. Paulos, M. Yim, A. Ribeiro, and V. Kumar, "Coverage control in multi-robot systems via graph neural networks," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 8787–8793.
- [74] M. Tzes, N. Bousias, E. Chatzipantazis, and G. J. Pappas, "Graph neural networks for multi-robot active information acquisition," *arXiv preprint arXiv:2209.12091*, 2022.
- [75] F. Chen, J. Wang, T. Shan, and B. Englot, "Autonomous exploration under uncertainty via graph convolutional networks," in *Proceedings of the International Symposium on Robotics Research*, 2019.
- [76] A. Agarwal, S. Kumar, K. Sycara, and M. Lewis, "Learning transferable cooperative behavior in multi-agent teams," in *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, ser. AAMAS '20. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2020, p. 1741–1743.
- [77] H. Zhang, J. Cheng, L. Zhang, Y. Li, and W. Zhang, "H2gnn: Hierarchical-hops graph neural networks for multi-robot exploration in unknown environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3435–3442, 2022.
- [78] E. Tolstaya, J. Paulos, V. Kumar, and A. Ribeiro, "Multi-robot coverage and exploration using spatial graph neural networks," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8944–8950.