# Research Statement

FANG, Yuan
School of Computing and Information Systems, Singapore Management University
Tel: (65) 6808-5150; Email: yfang@smu.edu.sg
14 December 2024

## I. Background and Overview

Graphs are prevalent in real-world applications, for they can model not only individual data entities, but also interactions between these entities. Example graphs include social networks, transportation and communication systems, financial networks, citation networks, and biochemical graphs [WPC20]. To gain insights into such data, my research (Figure 1) has undertaken *graph learning and mining*. I particularly focus on three topics: a) learning graph representations, B) multi-modal graph learning, and C) data efficiency for graph learning, as well as their various applications.

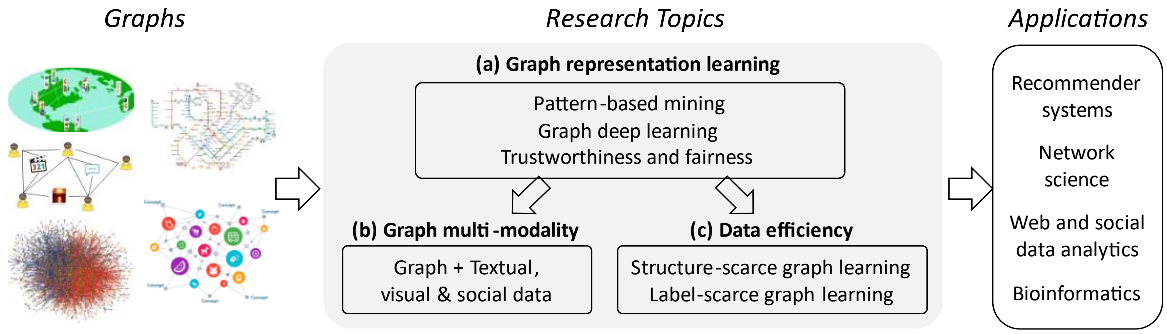


**Figure 1**: Overall research theme – Graph learning and mining.

Overall, my research is driven by the need to develop principled methodology for various fundamental problems on graphs. My work on graph representations investigates the underlying principles and mechanisms in graph-structured data, which builds a mathematical and algorithmic foundation for exploring different aspects and applications on graphs. On top of the fundamental graph representation learning, my choice of specific research problems has been motivated by opportunities and challenges in graph learning and mining. On one hand, given the variety and volume of multi-modal data associated with graph structures, there is an urgency to leverage graph multi-modality to augment traditional problems on graphs. Hence, I have explored the setting of graphs augmented by textual, visual and social data. These data often complement the graph structures to enable more effective and robust learning. On the other hand, one of the biggest hurdles faced by the deep learning community is data efficiency, particularly the lack of labelled data. The same problem persists in graph learning, where we need to address two aspects of data efficiency: when the structures are scarce, and when the labels are scarce. Addressing these research problems is crucial to not only the understanding of fundamental principles on graphs, but also the enabling of practical applications on graphs.

## II. Summary and Evolution of My Research

As new opportunities and challenges arise, my research has evolved around several important topics in graph learning and mining (Table 1).

My early research has studied the fundamental principles of graph representations based on structural patterns, ranging from simple links [VLDB13] to more complex semantic structures, most notably the concept of *metagraph* [ICDE16, TKDE19]. Metagraph is a representative semantic structure, which have since been widely used in the community for various applications and problems on graphs [ZYL17, SZC19], including in some of my own research [MET17, TKDE20, CIKM21]. Since I started my career as an assistant professor, my research on graph representations has gradually taken on deep learning-based approaches, given the superior performance of neural networks. Nonetheless, such more recent work [KDD19, TKDE20, IJCAI21, AAAI23a] are built upon the insights of pattern-based representations. In addition to performance-oriented methods, in recent years I have also explored the trustworthiness and fairness in representation learning [TOIS23, AAAI23b, NeurIPS23a]. This line of work has become increasingly important, as it addresses critical concerns about the ethical and societal impacts of artificial intelligence [KUR22].

Building upon graph representation learning, my research has also explored graph multi-modality, given the opportunities presented in the big data era. Specifically, abundant multi-modal data, including textual, visual and social data, are often associated with graphs. These data encompass rich information to complement graph structures and can be immensely valuable to enhancing graph learning models for various graph-based applications. While my earlier works [TACL14, IJCAI17] mostly leverage the structural patterns, a more recent study [SIGIR23] have leveraged graph neural networks as the backbone.

While deep learning has been widely adopted in different communities including graph learning, its success critically depends on the availability of large-scale labelled data. This limitation has motivated research on data-efficient learning, in which we aim to learn a strong model even without large-scale data. Hence, my more recent research has paid significant attention to the data efficiency problem on graphs, in which there is not only the label-scarcity issue [AAAI21b, SIGIR21], but also the structure-scarce issue [CIKM20, KDD20, KDD21]. Recently, graph pre-training on large-scale label-free graphs has emerged as a promising direction [AAA21a, CIKM21, NeurIPS23b]. Inspired by prompt-based learning in large language models, I have furthered our graph pre-training research into prompt learning on graphs [WWW23, CIKM23, AAAI24]. These studies have established a parameter-efficient adaptation for pre-trained graph models, showcasing great generalization ability across different downstream tasks with limited task-specific labels.

| Topics | Subtopics | Period | Selected publications |
|---|---|---|---|
| (a) Graph representation learning | Pattern-based mining | 2010–2019 | VLDB13, ICDE16, MET17, TKDE19 |
| | Deep graph learning | 2018–present | KDD19, TKDE20, IJCAI21, WWW22, AAAI23a, WWW24 |
| | Trustworthiness & fairness | 2022–present | TOIS23, AAAI23b, NeurIPS23a |
| (b) Graph multi-modality learning | | 2013–present | TACL14, IJCAI17, SIGIR23 |
| (c) Data-efficient graph learning | Structure-scarce graphs | 2019–2023 | KDD20, CIKM20, KDD21, TKDE23 |
| | Label-scarce graphs | 2020–present | AAAI21a, AAAI21b, SIGIR21, CIKM21, WWW23, CIKM23, NeurIPS23b, AAAI24 |

**Table 1**: Topical and chronological evolution of my research.

## III. Notable Recognition of Research

**Research output.** I have published over 70 conference and journal papers in leading data mining, machine learning and the broader AI venues such as NeurIPS, KDD, WWW, SIGIR, AAAI and TKDE. My publications have garnered a total citation of 2780+ and h-index of 31 on Google Scholar (see Appendices 1 & 2 for more details). Most notably, our recent work on GraphPrompt [WWW23] has been recognized among the Most Influential WWW Papers (ranked #5 in WWW'23) by Paper Digest in May 2024, and our work on PageRank [VLDB13] has been featured in the special issue on Best Papers of VLDB'13 (only 4 selected out of 127).

**Fundings and grants.** My research activities were or are supported in part by Singapore's Ministry of Education (MOE Tier 2), A*STAR, AI Singapore, Alibaba Group and DBS Bank, as well as SMU internal funding. Including the period before joining SMU, I served or am serving as the PI, Co-PI or project-level PI across 8 projects. Since I joined SMU in July 2018, as the PI or project-level PI, I have secured a total external funding size of SGD 1.6 million (~USD 1.2 million), and managed and executed externally funded projects worth SGD 2.1 million in total.

**Academic supervision.** In conjunction with my research, I have supervised PhD students and research staff who have become promising young researchers. Two of my PhD students have graduated, namely, Zhihao Wen (2023) and Zhongzhou Liu (2024), who joined Huawei's research lab in Singapore. Both of them were awarded the prestigious SMU Presidential Doctoral Fellowship and made the Dean's List, and Zhongzhou also received the NeurIPS 2023 Scholar Award. My former postdoctoral fellow, Zemin Liu, joined Zhejiang University as an Assistant Professor in May 2024. My former research assistant, Jiaqi Shi, joined the University at Buffalo as an Assistant Professor in Fall 2024.

**Professional activities.** I was an invited speaker at the Social Networks Analysis Forum, 8th China Conference on Data Mining in 2020 on my series of work on metagraphs, a keynote speaker in prestigious workshops collated at WSDM'23 and RecSys'23 on my work in data-efficient graph learning, and an invited panellist in the WWW'23 Workshop on Graph Foundation Models. I also served as an Area Chair on the program committee of WWW'24, a Workshop Co-Chair for WSDM'23, SPC/PC members and Session Chairs of top conferences, and reviewers for leading journals. For professional societies, I have been the Secretary of Singapore ACM SIGKDD Chapter since 2022 and served as the PC Co-Chair for the local chapter's Symposium in 2023. I have also been elevated to a Senior Member of IEEE in 2024.

## IV. Key Research Findings by Area

### (a) Graph representation learning

Learning effective graph representations has been and remains a core problem for graph-based tasks.

**Pattern-based approaches.** My earlier research revolved around directly utilizing structural patterns in graphs. In real-world applications, entities are often interlinked to form heterogeneous graphs, where different semantics exist between nodes. For instance, on a social network, various relationships between users exist: classmates, family, colleagues, etc. We proposed *metagraph representations* [ICDE16] as a novel concept to characterise these different semantic relationships, which have shown strong results in many graph-based applications such as social recommendation [ICDE16] and bioinformatics [MET17]. Taking a step further, we have also explored metagraphs as a universal form of node and edge representations [TKDE19], demonstrating its superior performance in more downstream tasks. Apart from using metagraphs as explicit representations, they can also serve as a tool for graph deep learning [TKDE20] and pre-training [CIKM21].

**Graph deep learning.** I have also investigated various techniques of graph embedding and neural networks for representation learning on different kinds of graphs. On general graphs, we have studied the node-wise adaptation [IJCAI21] and edge-centric message-passing [AAAI23a] for graph neural networks. On heterogenous graphs, our research has investigated adversarial learning [KDD19] and metagraph-guided embedding [TKDE20]. On dynamic graphs, I have studied the Hawkes process on graph neural networks [WWW22] and the Transformer-based approach [WWW24]. Due to the ability to fit complex, nonlinear functions, deep graph representation learning often achieves state-of-the-art performance in various domains such as bioinformatics and recommender systems.

**Trustworthiness and fairness.** Traditionally, graph representation learning focuses on improving performance in graph-based applications but pays little attention to the trustworthiness and fairness of the model predictions. My recent work on structural fairness in graph neural networks [AAAI23b] aims to develop methods for learning fair representations of nodes, particularly focusing on addressing the disparities in structural resources (e.g., social capital) among these nodes. The goal is to ensure that the algorithmic outcomes are equitable, regardless of the varied structural resources of the nodes. The fairness concept can also be extended to user-item bipartite graphs in recommender systems. On this bipartite graph, we investigated the issue of popularity bias [TOIS23], aiming to achieve fairness for both users and sellers, regardless of their existing popularity (Figure 2). On recommender systems, we also explored propensity estimation for causality-based recommender systems [NeurIPS23a]. This study enables a more practical setup for causality-based recommendations that are not only effective but also more trustworthy and explainable, by producing outcomes grounded in causal effect rather than mere correlations in historical data.

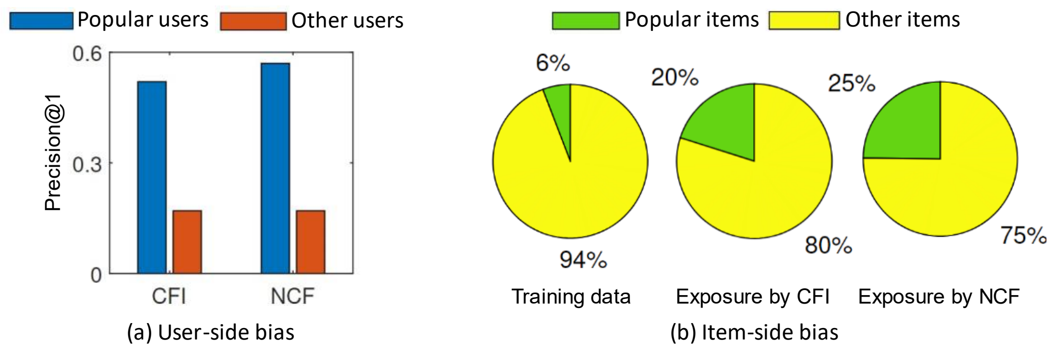


**Figure 2**: Popularity bias in conventional recommender models, e.g., CFI [LZL18] and NCF [HLZ17]. (a) User-side bias: popular users enjoy much higher recommendation accuracy than others; (b) Item-side bias: popular items receive much more exposure (20–25%) than their proportion in the training data (6%).

## (b) Graph multi-modality learning

Many problem statements often involve other kinds of data in addition to explicit graph structures, including visual, textual, and social data. These data either enable us to construct new graphs, or to complement existing graphs to improve learning or to enable new tasks. We refer to such research as multi-modal graph learning, with the goal of exploiting additional knowledge from one or more complementary modalities beyond graph structures.

In an earlier work [IJCAI17], we exploited knowledge graphs to improve the object detection task in images. Specifically, knowledge graphs contain commonsense knowledge that can relate different objects in an image, e.g., pets like cats and furniture like dining table often appear together in households. Such knowledge would improve detection recalls in home scenes: the detections of pet and furniture mutually reinforce each other, should one of them has low initial confidence. Besides visual data, textual information, such as abstracts for papers on a citation graph (Figure 3), or item descriptions in a recommender system, can enrich the interactions between nodes. Inspired by the success of prompt-based learning in pre-trained language models, we investigated prompt tuning techniques for such text-attributed graphs [SIGIR23], achieving state-of-the-art performance in node classification tasks, significantly surpassing graph-only or text-only models.

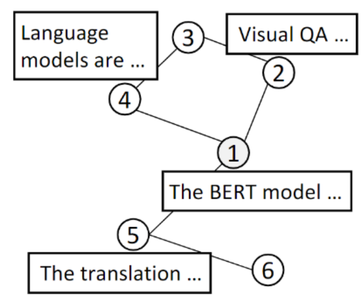


**Figure 3**: An example text-attributed citation graph, where each node represents a paper with a textual description (e.g. title or abstract).

However, our approach on text-attributed graphs [SIGIR23] could not work if there is no graph structure relating the texts. An alternative is to use social meta-data to construct auxiliary graphs. In our entity linking study [TACL14], we constructed spatial and temporal graphs for entities appearing on Twitter (known as X now), so that entities close to each other in either space or time can be connected based on the meta-data of the tweets (i.e., timestamps and geotags). The connections reveal the relatedness between entities, which proves beneficial to the entity linking task.

## (c) Data efficiency in graph learning

Learning with data efficiency has gained significant traction in recent years as a means to overcome the requirement of large-scale labelled data in conventional supervised deep learning. We have explored two major aspects of data efficiency in graph learning.

**Structure-scarce graph learning.** On many graphs, we observe that there is often a long tail of nodes with very few links. In general, the node degrees vary considerably across the network and are not uniformly distributed (Figure 4a). Learning representations for these structure-scarce “tail” nodes are more challenging than nodes with rich structures (Figure 4b), presenting a novel problem that is often neglected by conventional graph learning. Leveraging the power of meta-learning, we formulated the problem as a few-shot regression task and proposed meta-tail2vec [CIKM20], a first attempt on this problem. However, metatail2vec is a two-stage method that improves the tail node embedding through a post-processing step. Thus, we further proposed an end-to-end tail node representation learning framework for graph neural networks [KDD21]. Similarly, the cold-start recommendation problem also suffers from the scarcity of structures connecting new users and items. Thus, we formulated the cold-start problem as a few-shot link prediction task [KDD20] and addressed it from the model and data levels with a co-adaptation meta-learner.

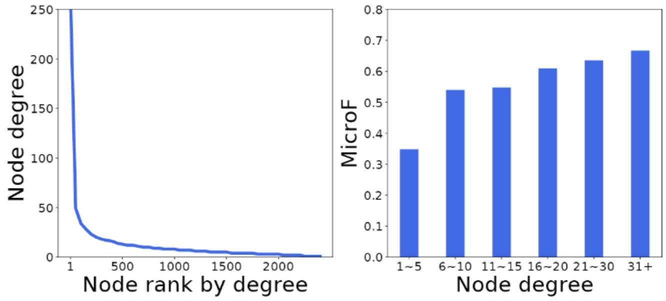

(a) Degree distribution (b) Classification performance

**Figure 4**: Relationship between node degrees and the quality of learned representations on a typical graph.

**Label-scarce graph learning.** Similar to other supervised models, state-of-the-art graph neural networks depend on abundant labelled data to achieve optimal performance. However, in real-world applications, many tasks often lack abundant labelled data. One common scenario is the few-shot node classification on graphs, in which some novel classes only have one or few examples. For instance, on a citation graph, while Markov chains is a mature topic with many labelled examples, algorithmic explainability and fairness is relatively new with few labelled examples. To address few-shot learning on graphs, we resort to the framework of meta-learning, while simultaneously exploiting graph-specific characteristics including the long-range dependencies between nodes, and the global graph contexts [AAAI21b]. We have also explored inductive graph learning across graphs, where the trained model on an existing set of graphs can be transferred to new graphs in the same feature space, reducing the need for labels on the new graphs [SIGIR21].

Meanwhile, to make use of the vast availability of “label-free” graphs (i.e., graphs without any label for downstream tasks), pre-training has emerged as a promising direction to capture inherent graph properties in a task-agnostic manner, which can be transferred to different downstream tasks. To capture such properties, we designed various pre-training objectives [AAAI21a, CIKM21, NeurIPS23b]. Conventionally, the pre-trained graph models can be adapted to diverse downstream tasks through fine-tuning using task-specific labels. However, full fine-tuning is not only inefficient, but also suboptimal especially when there are few task-specific labels due to the divergence between pre-training and downstream objectives. To address the objective differences, we draw inspiration from prompt-based learning in pre-trained language models, in which prompts are designed to generalize the pre-trained model to a wide range of downstream tasks using limited task-specific labels. Consequently, we designed GraphPrompt [WWW23], one of the pioneering studies that attempt to unify pre-training and downstream tasks in a prompt-tuning framework for graph learning. We further applied prompt-based learning to address anomaly detection in e-commerce [CIKM23] and handle heterogeneous graphs [AAAI24].

## V. Future Research Agenda

My future work will still be anchored on graph learning and mining, as many research questions and opportunities remain open in the three major topics of my research (Figure 1).

**Towards trustworthy and fair graph learning.** While we have conducted some preliminary work in trustworthiness [NeurIPS23a] and fairness [TOIS23, AAAI23b], both our research and the broader literature in the community are still in their early stages, with many challenges remaining.

First, different perspectives of fairness and various forms of biases exist. For example, fairness can be defined based on statistical parity and equal opportunities, while various biases such as self-selection, conformity, exposure, and popularity are prevalent in recommender systems [CDW23]. These definitions may sometimes conflict with each other, driving the outcomes in opposing directions. Thus, it is crucial to design an adaptive framework that can accommodate different definitions of fairness and bias on graphs in a coherent manner.

Second, how do we study various related concepts in trustworthy graph learning systematically? For example, fairness and explainability are often not independent but have intricate relationships. A decision that can be explicitly explained often helps mitigate perceived unfairness. Thus, considering two or more aspects jointly could improve the trustworthiness of graph learning.

**Towards multi-modal graph foundation models.** Recently, large language models (LLMs) or foundation models [BHA21] have unified the pre-training of a broad range of language-based data across many domains, enabling an exciting array of downstream tasks. Given their successes, it is intriguing to ask: *Is it possible to develop a graph foundation model [LYL23] on broad graph data across diverse domains, which can be subsequently used to address a wide range of downstream tasks on graphs?*

The first major hurdle towards this goal lies in the challenge that graphs from different domains often exhibit divergent characteristics. In LLMs, textual contents from multiple domains are inherently connected through a common set of natural language tokens (i.e., words). In contrast, for multi-domain graphs, the tokens (i.e., nodes of the graphs) can have little overlap across graphs in different domains, ranging from users on a social network, to corporations in a financial network and atoms in a molecular graph. As a result, directly applying existing graph pre-training techniques to multi-domain graphs often leads to domain conflicts, as our preliminary study shows (Figure A5). Thus, it is imperative to study pre-training methodologies for multi-domain graphs. Meanwhile, the domains must be carefully scoped *within the same discipline* to ensure there are some common tokens across domains. For example, protein-protein interactions, gene regulatory networks, and metabolic pathways are multi-domain graphs within the biomedical discipline, while customer relationships and transaction networks are multi-domain graphs within the finance discipline. In each discipline, there could be a set of explicit or latent common tokens to bridge the multi-domain graphs for synergistic pre-training.

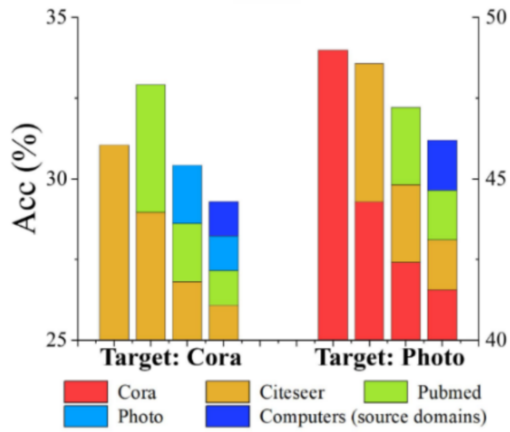


**Figure A5**: Accuracy of one-shot node classification tasks on two target domains, as more source domains are added to pre-training using a popular graph pre-training approach [VFH19]. Pre-training with more source domains often results in reduced accuracy, indicating domain conflicts.

Second, how can we integrate large-scale visual and textual data to complement graph learning? Our existing research on graph multi-modal learning [TACL14, IJCAI17, SIGIR23] has shed light on this possibility, albeit on a smaller scale. Towards multi-modal graph foundation models, it would be crucial to explore the integration of graph structures with existing large language and vision models. One plausible starting point is our research on text-attributed graphs [SIGIR23], where we align the graph and language models with each other. However, considering the scale of the LLMs, injecting the smaller graph models into existing LLMs presents a more feasible direction.

Third, I plan to investigate parameter-efficient fine-tuning techniques for pre-trained graph models beyond our prompt-based work [WWW23], such as adapter learning and low-rank adaptation [ACL24]. These techniques would further enhance the generalization ability across different downstream tasks in a parameter-efficient manner, potentially contributing to the development of graph foundation models, driving future research on offering universal transferability and adaptability to broad graph data and tasks.

## VI. References

### (a) References to my work (selected)

[VLDB13] F Zhu, **Y Fang**, K CC Chang, J Ying. Incremental and Accuracy-Aware Personalized PageRank through Scheduled Approximation. In *VLDB* 2013.

[TACL14] **Y Fang**, MW Chang. Entity Linking on Microblogs with Spatial and Temporal Signals. *TACL* 2(Oct), 2014.

[ICDE16] **Y Fang**, W Lin, VW Zheng, M Wu, K CC Chang, XL Li. Semantic Proximity Search on Graphs with Metagraph-based Learning. In *ICDE* 2016.

[IJCAI17] **Y Fang**, K Kuan, J Lin, C Tan, V Chandrasekhar. Object Detection Meets Knowledge Graphs. In *IJCAI* 2017.

[MET17] S Kircali, **Y Fang**, M Wu, X Xiao, XL Li. Disease Gene Classification with Metagraph Representations. *Methods* 131, 2017.

[KDD19] B Hu, **Y Fang**, C Shi. Adversarial Learning on Heterogeneous Information Networks. In *KDD* 2019.

[TKDE19] **Y Fang**, W Lin, VW Zheng, M Wu, J Shi, K CC Chang, XL Li. Metagraph-based Learning on Heterogeneous Graphs. *IEEE TKDE* 33(1), 2019.

[TKDE20] W Zhang, **Y Fang**, Z Liu, M Wu, X Zhang. mg2vec: Learning Relationship-Preserving Heterogeneous Graph Representations via Metagraph Embedding. *IEEE TKDE* 34(3), 2020.

[KDD20] Y Lu, **Y Fang**, C Shi. Meta-learning on Heterogeneous Information Networks for Cold-start Recommendation. In *KDD* 2020.

[CIKM20] Z Liu, W Zhang, **Y Fang**, X Zhang, SCH Hoi. Towards Locality-Aware Meta-Learning of Tail Node Embeddings on Networks. In *CIKM* 2020.

[AAAI21a] Y Lu, X Jiang, **Y Fang**, C Shi. Learning to Pre-train Graph Neural Networks. In *AAAI* 2021.

[AAAI21b] Z Liu, **Y Fang**, C Liu, SCH. Hoi. Relative and Absolute Location Embedding for Few-Shot Node Classification on Graph. In *AAAI* 2021.

[KDD21] Z Liu, K Nguyen, **Y Fang**. Tail-GNN: Tail-Node Graph Neural Networks. In *KDD* 2021.

[SIGIR21] Z Wen, **Y Fang**, Z Liu. Meta-Inductive Node Classification across Graphs. In *SIGIR* 2021.

[CIKM21] X Jiang, Y Lu, **Y Fang**, C Shi. Contrastive Pre-training of GNNs on Heterogeneous Graphs. In *CIKM* 2021.

[IJCAI21] Z Liu, **Y Fang**, C Liu, SCH. Hoi. Node-wise Localization of Graph Neural Networks. In *IJCAI* 2021.

[WWW22] Z Wen, **Y Fang**. TREND: TempoRal Event and Node Dynamics for Graph Representation Learning. In *WWW* 2022.

[WWW23] Z Liu, X Yu, **Y Fang**, X Zhang. GraphPrompt: Unifying Pre-Training and Downstream Tasks for Graph Neural Networks. In *WWW* 2023.

[AAAI23a] X Yu, Z Liu, **Y Fang**, X Zhang. Learning to Count Isomorphisms with Graph Neural Networks. In *AAAI* 2023.

[AAAI23b] Z Liu, T Nguyen, **Y Fang**. On Generalized Degree Fairness in Graph Neural Networks. In *AAAI* 2023.

[TOIS23] Z Liu, **Y Fang**, M Wu. Mitigating Popularity Bias for Users and Items with Fairness-centric Adaptive Recommendation. In *ACM TOIS* 41(3), 2023.

[TKDE23] Z Liu, **Y Fang**, W Zhang, X Zhang, SCH Hoi. Locality-Aware Tail Node Embeddings on Homogeneous and Heterogeneous Networks. In *IEEE TKDE* 36(6), 2023.

[NeurIPS23a] Z Liu, **Y Fang**, M Wu. Estimating Propensity for Causality-based Recommendation without Exposure Data. In *NeurIPS* 2023.

[NeurIPS23b] D Bo, **Y Fang**, Y Liu, C Shi. Graph Contrastive Learning with Stable and Scalable Spectral Encoding. In *NeurIPS* 2023.

[SIGIR23] Z Wen, **Y Fang**. Augmenting Low-Resource Text Classification with Graph-Grounded Pre-training and Prompting. In *SIGIR* 2023.

[CIKM23] Z Wen, **Y Fang**, Y Liu, Y Guo, S Hao. Voucher Abuse Detection with Prompt-based Fine-tuning on Graph Neural Networks. In *CIKM* 2023 (Applied Research).

[LYL23] J Liu, C Yang, Z Lu, et al, **Y Fang**, et al. Towards Graph Foundation Models: A Survey and Beyond. In *arXiv*:2310.11829.

[AAAI24] X Yu, **Y Fang**, Z Liu, X Zhang. HGPrompt: Bridging Homogeneous and Heterogeneous Graphs for Few-shot Prompt Learning. In *AAAI* 2024.

[WWW24] Y Wu, **Y Fang**, L Liao. On the Feasibility of Simple Transformer for Dynamic Graph Modeling. In *WWW* 2024.

[ACL24] Z Wen, J Zhang, **Y Fang**. SIBO: A Simple Booster for Parameter-Efficient Fine-Tuning. In *ACL* (Findings) 2024.

## (b) Other references

[ZYL17] H Zhao, Q Yao, J Li, Y Song, DL Lee. Meta-Graph Based Recommendation Fusion over Heterogeneous Information Networks. In *KDD* 2017.

[HLZ17] Neural collaborative filtering. X He, L Liao, H Zhang, L Nie, X Hu, TS Chua. In *WWW* 2017.

[LZL18] Y Liu, L Zhao, G Liu, X Lu, P Gao, XL Li, Z Jin. Dynamic Bayesian Logistic Matrix Factorization for Recommendation with Implicit Feedback. In *IJCAI* 2018.

[VFH19] Petar Veličković, William Fedus, William L. Hamilton, et al. Deep Graph Infomax. In *ICLR* 2019.

[SZC19] A Sankar, X Zhang, KCC Chang. Meta-GNN: Metagraph Neural Network for Semi-Supervised Learning in Attributed Heterogeneous Information Networks. In *ASONAM* 2019.

[WPC20] Z Wu, S Pan, F Chen, G Long, C Zhang, PS Yu. A Comprehensive Survey on Graph Neural Networks. In *IEEE TNNLS* 32(1), 2020.

[BHA21] R Bommasani, DA Hudson, E Adeli, et al. On the Opportunities, Risks of Foundation Models. In *arXiv*:2108.07258.

[KUR22] D Kaur, S Uslu, KJ Rittichier, A Durresi. Trustworthy Artificial Intelligence: a Review. In *CSUR* 55(2), 2022.

[CDW23] Bias and Debias in Recommender System: A Survey and Future Directions. J Chen, H Dong, X Wang, et al. In *ACM TOIS* 41(3), 2023.