Research Statement

Dr. Yuan FANG
School of Computing and Information Systems
Singapore Management University
Email: yfang@smu.edu.sg

25 December 2023

I. Background & Overview

Graphs are prevalent in real-world datasets, for they can model not only individual data entities, but also interactions between these entities. Example graphs include the Web, social networks, transportation and telecommunication systems, scholarly citation networks, and protein interaction networks. To gain insights into such data, my research (Figure 1) has undertaken *learning and mining on graphs*. In particular, I focus on three sub-areas: A) designing and learning graph representations, B) multi-modal graph-based learning, and C) data efficiency and scalability on graphs, as well as their various applications.

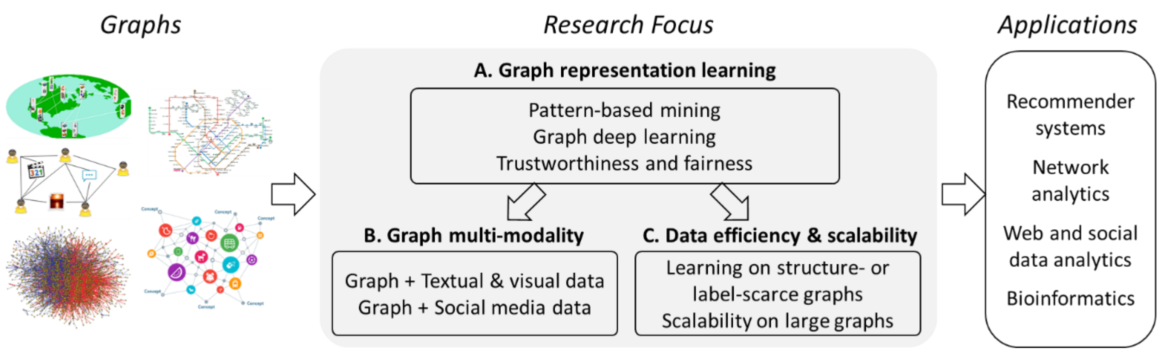

Figure 1: Overall research theme – Learning and mining on graphs.

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. In a finer granularity, my choice of specific research problems has been motivated by opportunities and challenges in the learning and mining of graph data. On the 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 learning and mining on graphs. Hence, I have explored the setting of "graph + X", where X can be textual and visual data, or social media data. These data often complement the graph structures to enable more effective and robust learning. On the other hand, one of the biggest hurdle faced by the deep learning community is data efficiency in the absence of abundant labelled data. The same problem persists in learning on graphs. Moreover, in graph 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.

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

My early research has studied the fundamental principles of graph representations based on structural patterns, ranging from simple links to more complex semantic structures. Compared to links, semantic structures embody higher-order patterns and thus can capture complex heterogeneous relationships on graphs. Hence, semantic structures provide a higher capacity for graph representations. In particular, my research proposed metagraph [ICDE16b], a representative semantic structure which have since been widely used in the community for various applications and problems on graphs, including in my own research [Method17, TKDE19, CIKM21b]. Since I joined SMU, 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, TKDE21b, IJCAI21, WWW23b, AAAI23a] are built upon the insights and foundation of pattern-based representations. In addition to the performance-oriented evaluation, in recent years I have also explored the trustworthiness and fairness in representation learning [TOIS23, AAAI23b, NeurIPS23a]. This line of work has become increasingly, as it addresses critical concerns about the ethical and societal impacts of artificial intelligence.

Besides investigating the fundamental principles and methodologies on graphs, my research has also explored graph multi-modality, given the opportunities presented in the big data era. Specifically, the abundance of multi-modal data associated with graphs, including textual and visual data, and social media data. 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 earlier works [TACL14, IJCAI17] mostly leverage the foundational structural patterns, more recent works [CIKM20a, CIKM20b, ECMLPKDD20b, SIGIR23] have leveraged graph embedding and neural networks.

While deep learning has been widely adopted in different communities including learning on graphs, 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 good 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 label-scarcity issue [AAAI21b, SIGIR21], but also structure-scarce issue [BIOINF20, CIKM20b, KDD20, KDD21a, WWW22b]. Moreover, making use of large-scale label-free graphs in a pre-training setting is also a promising direction, as demonstrated in a series of work [AAA21a, CIKM21b, KDD21b, NeurIPS23b]. More recently, as inspired by prompt-based learning in pre-trained language models, we furthered our graph pre-training research into prompt learning on graphs [WWW23a, CIKM23, AAAI24]. Similar to prompt on pre-trained language models, prompt on pre-trained graph models show great generalization ability across different downstream tasks with very limited task-specific supervision. On a separate line, in order to handle massive graphs (e.g., the label-free graphs used in pre-training), some of my research also studied the scalability of graph algorithms, ranging from pattern-based algorithms [VLDB13, VLDBJ15, TKDE21a] to deep learning-based algorithms [TKDD20, KDD21b].

<table>
<tr><th>Theme</th><th>Detailed topic</th><th>Research publications (Chronologically ordered)</th></tr>
<tr><td rowspan="3">Graph representation learning</td><td>Pattern-based mining</td><td>WSDM11, SIGIR12, ICDE13, ICML14, ICDE16a, ICDE16b, Method17, TKDE19</td></tr>
<tr><td>Deep graph learning</td><td>BMC18, ICDM18, KDD19, IPM20, SDM20, TKDE20, IJCAI21, TKDD21, JBHI21, TKDE21b, WWW22a, WWW23b, AAAI23a, TKDE23</td></tr>
<tr><td>Trustworthiness & fairness</td><td>TOIS23, AAAI23b, NeurIPS23a</td></tr>
<tr><td rowspan="2">Graph multi-modality</td><td>Graph + text & visual</td><td>IJCAI17, CIKM20a, CIKM21b, SIGIR23</td></tr>
<tr><td>Graph + social data</td><td>TACL14, ECMLPKDD20b</td></tr>
<tr><td rowspan="3">Data efficiency & scalability in graph learning</td><td>Label-scarce graphs</td><td>AAAI21a, AAAI21b, SIGIR21, CIKM21b, IJCAI22, BIOINF22, WWW23a, CIKM23, NeurIPS23b, AAAI24</td></tr>
<tr><td>Structure-scarce graphs</td><td>BIOINF20, KDD20, CIKM20b, KDD21a, WWW22b</td></tr>
<tr><td>Scalability</td><td>VLDB13, VLDBJ15, TKDD20, TKDE21a, KDD21b</td></tr>
</table>

Table 1: Topical and chronological evolution of my research.

Notable recognition of research. My research activities were or have been supported in part by Ministry of Education Singapore (MOE Tier 2), A*STAR, AI Singapore, Alibaba Group and DBS Bank, as well as SMU internal funding (MOE Tier 1 and LARC funding). I have served as the PI or Co-PI/Project-level PI across 7 projects. Since I joined SMU in July 2018, I have managed a total funding size of SGD 2.1 million from externally funded projects as the PI or project-level PI. I have published over 60 conference and journal papers in leading data mining, machine learning and the broader AI venues such as NeurIPS, KDD, WWW, SIGIR, ICML, AAAI and TKDE. As of the time of writing, my publications have attracted a total citation of

2200+ and h-index of 28 according to Google Scholar. I was an invited speaker at the Social Networks Analysis Forum, 8th China Conference on Data Mining in 2020, on my series of work exploiting high-order semantic structures in heterogeneous information networks, and a keynote speaker in prestigious workshops collated at WSDM'23 and RecSys'23 on my work in data-efficient graph learning. I also served as an Area Chair on the program committee of WWW'24. My research [VLDBJ15] was featured in the "**Best Papers of VLDB13**" special issue of the VLDB Journal, which is a recognition and extended version of our work on efficient personalized PageRank computation [VLDB13].

II. Key Research Findings by Theme

A. Graph representation learning

Structural pattern-based approaches. My earlier research revolved around directly utilizing structural patterns in graphs. Unlike traditional flat data, graph structures explain complex interactions between data entities, and thus are crucial towards data-driven tasks. Leveraging on the simple link structures, we investigated semi-supervised learning on graphs [ICML14]. The resulting graph-based probabilistic framework unifies the underlying principle in our previous random walk models [WSDM11, ICDE13]. We further considered heterogeneous link structures [SIGIR12], as well as extended the learning objective on individual nodes to a set of nodes [ICDE16a]. In summary, link-based learning on graphs enable us to improve various tasks on graphs, including node classification and ranking, information extraction and data-driven crawling.

I have also studied higher-order semantic structures beyond simple link structures. In real-world scenarios, objects are often interlinked to form heterogeneous graphs, where different semantics exist between nodes. For instance, the below social network (Figure 2) contains users of different semantic relationships: some are classmates, some are family, and some are colleagues. The multitude of semantics arises from various types of nodes and their different interactions. We have proposed metagraph representations [ICDE16b] as a novel means to characterise these different semantic classes, which have shown very promising results in our studies on proximity ranking [ICDE16b] and node classification [Methods17]. 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 including clustering and relationship prediction. Apart from using metagraphs as explicit representations, they can also serve as the foundation in graph deep learning [TKDE20] and graph pre-training [CIKM21b] on heterogeneous graphs.

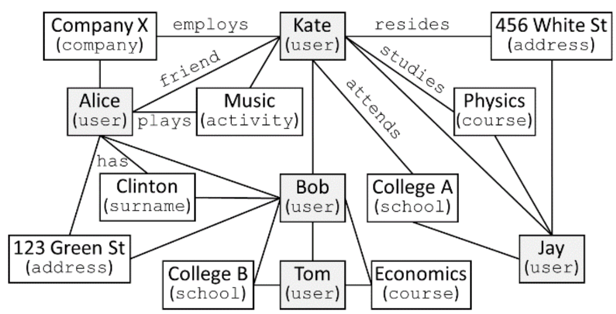

Figure 2. Example social network with rich semantics.

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 of graph neural networks [IJCAI21], edge-centric message-passing [AAAI23a], and latent heterogeneous graph neural networks [WWW23b]. On heterogenous graphs, we have investigated neighborhood propagation [ICDM18], adversarial learning [KDD19] and metagraph-guided embedding [TKDE20]. On dynamic graphs, we have researched the neural attention mechanism [ECMLPKDD20a] and Hawkes process [ECMLPKDD21, WWW22a]. On attributed graphs, we have explored the integration of structures and attributes [IPM20, SIGIR20]. In summary, these works all leverage artificial neural networks for graph representation learning. Due to the ability to fit complex, nonlinear functions, neural networks-based graph representation learning often achieve state-of-the-art performance in various domains such as bioinformatics [BMC18, JBHI21] and recommendation systems [SDM20, ECMLPKDD20a, TKDE23].

Trustworthiness and fairness. Traditionally, graph representation learning primarily focuses on improving performance in graph-based applications, often emphasizing accuracy, scalability, and efficiency. This focus typically involves developing algorithms to effectively capture the structure and features of graphs, but pays little attention to other crucial aspects, such as the trustworthiness and fairness of the outcomes delivered by the models, which entail profound ethical and societal consequences. More specifically, our 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 recommender systems, which deals with a user-item bipartite graph. On this bipartite graph, we investigate the issue of popularity bias [TOIS23], aiming to achieve fairness on both user and item side, regardless of their existing popularity. On recommender systems, we also explore propensity estimation for causality-based recommendation systems [NeurIPS23a]. This study focuses on developing recommendation algorithms that are not only effective but also more trustworthy and explainable, by producing outcomes that are grounded in causal effect rather than mere correlations.

B. Graph multi-modality learning

Many problem statements often involve other kinds of data in addition to explicit graph structures, including visual and textual data and social meta-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 such research as multi-modal graph-based learning. Exploiting multi-modal data with graphs is a form of data enrichment to bring in knowledge that are not directly available from labeled data. There is a general consensus in the community that current machine learning approaches suffer from a significant knowledge gap. Additional knowledge from multi-modal data can potentially narrow the gap.

Graph + Textual & visual data. In my work on object detection in images [IJCAI17], we exploit knowledge graphs to improve the visual detection task (Figure 3). In particular, knowledge graphs contain commonsense knowledge that relate different objects in images. An example piece of commonsense knowledge is that pets (e.g., cats) and furniture (e.g., table) often appear together. 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. Alternatively, textual information, such as item descriptions in a recommendation scenario, can enrich the interactions between users and items to form a more dense graph structure, providing additional insights to boost recommendation performance [CIKM20a]. In another work, texts associated with graph nodes reveal multi-facet topical factors, which can guide finer-grained learning on graphs for both better model performance and interpretability [CIKM21b]. More recently, inspired by success of prompting in pre-trained language models, we have also investigated prompt tuning techniques for graph-grounded text classification [SIGIR23].

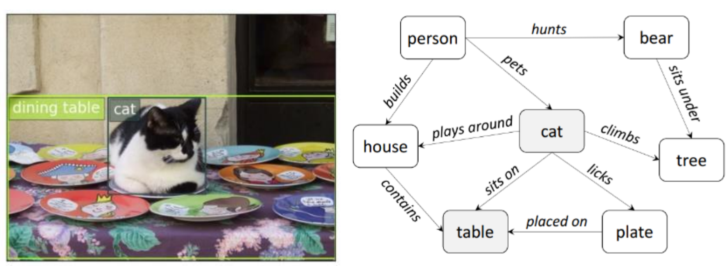

Figure 3. (left) An example task of object detection, to identify a dining table and a cat in the image. (right) Toy knowledge graph demonstrating the relationship between cats and tables.

Graph + Social data. In our entity linking study [TACL14], we construct spatial and temporal graphs for entities appearing on Twitter, so that entities that are close to each other in either space or time are connected based on the meta-data of the tweets (i.e., timestamp and geotagging of the tweet). The connections reveal the relatedness between entities, which proves beneficial to the entity linking task on Twitter. In a more recent work, we attempt to enrich collaborative filtering with a novel form of social meta-data known as "friend referral circle" [ECMLPKDD20b], where users are recommended with items liked or shared by their circle of friends. Leveraging the unique friend referral circle enables a more accurate modeling of social factors (e.g., user behaviors are more influenced by their friends who appears more authoritative), beyond just the homophily effect assumed in conventional social recommendation.

C. Data efficiency and scalability in graph learning

Learning with data efficiency has always been an important research topic, and has gained particular traction in recent years due to the rise of deep learning which often require large-scale, high-density data for optimal performance. To address the over-reliance on data, in particular on graphs, we have explored several directions of data efficiency. Besides, we have also studied the more conventional computational scalability problem on massive graphs.

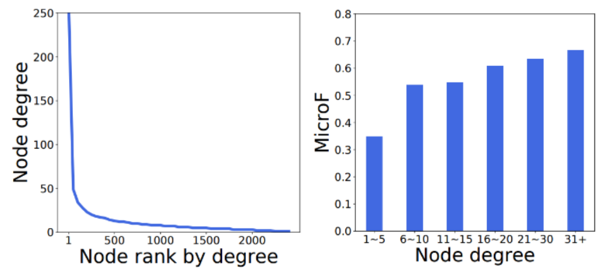

Figure 4. Relationship between node degree and the quality of learned representation on a typical graph. (left) Node degree distribution. (right) Classification performance w.r.t. node degrees.

Data efficiency on structure-scarce graphs. Structure-scarce graphs refer to graphs where the connective structures between nodes are very sparse. First, we investigate a dual dropout strategy on both nodes and edges for graph neural networks [BIOINF20], to increase the overall robustness of learning. Second, we observe that in a graph collection, there is often a long-tail of small graphs [WWW22b]. These tail graphs often lack distinguishable structures due to their small sizes. To help these tail graphs, we exploit a knowledge transfer based on pattern co-occurrence from head graphs, as the head graphs are larger in size and often have richer and more distinctive structures. Third, even if a graph is considered structure-rich on the whole, there still exist tail nodes with very few links. In other words, an individual tail node has very scarce structural connectivity, despite the abundance of links on other nodes. In general, the node degrees vary considerably across the network and are not uniformly distributed (Figure 4, left). The lack of structural connectivity on a tail node makes its representation more difficult to learn than nodes with abundant links (Figure 4, right). Representation learning for the tail nodes is thus a challenging and novel problem. Inspired by meta-learning, we formulate the problem as a few-shot regression task in our work meta-tail2vec [CIKM20b], a first attempt on this problem to our best knowledge. However, meta-tail2vec is a two-stage method that improves the tail node embedding through a post-processing step. Thus, we further propose an end-to-end tail node representation learning framework for graph neural networks [KDD21a]. Similarly, the cold-start recommendation problem also suffers from the scarcity of structures connecting new users and items. Thus, we formulate the cold-start problem as a few-shot link prediction task, and addressed it under the meta-learning paradigm as well [KDD20].

Data efficiency on label-scarce graphs. Like any other supervised machine learning models, state-of-the-art graph neural networks rely on a large number of labels for good performance. However, in reality, many tasks often lack abundant labeled data. One common scenario is the few-shot node classification task on a graph, in which some novel classes only have one or few examples. For instance, on a scholarly citation network, while Markov chains is a mature topic with many labeled examples, algorithmic explainability and fairness is relatively new with few labeled 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 [AAAI21a]. In another scenario, there might be out-of-distribution (OOD) classes on a graph, i.e., there is no label for these classes at all, and it is important to identify these OOD nodes for more accurate classification of in-distribution nodes [IJCAI22]. 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 given labels for downstream tasks), pre-training has been a promising direction to capture inherent graph properties in a task-agnostic manner, which can be transferred to different downstream tasks. Some of our works [KDD21b, CIKM21b, BIOINF22, NeurIPS23b] attempt to design various pre-training objectives to better capture such properties. The pre-trained graph models can be adapted to diverse downstream tasks through a fine-tuning step using some task-specific labels. However, pre-training is decoupled from fine-tuning, causing a divergence between their optimization objectives [AAAI21b]. To address the objective differences, we draw inspiration from prompt-based learning in natural language processing, in which prompts are designed to generalize the pre-trained model to a wide range of downstream tasks without the need for fine-tuning. Consequently, we have designed GraphPrompt [WWW23a], one of the pioneering works that attempt to unify pre-training and downstream tasks in a prompt-tuning framework for graph learning. We have further applied prompt-based learning to address anomaly detection in e-commerce [CIKM23], and handle heterogeneous graphs [AAAI24].

Scalability solutions on graphs. Finally, we explore scalable computational frameworks for massive graphs. Our work investigated fast approximation algorithms for computing Personalized PageRank [VLDB13, VLDBJ15] and SimRank [TKDE21a] on large graphs. The algorithm can speed up over existing methods by several folds with high accuracy and excellent scalability. More recently, as network embedding and graph neural networks emerge as the *de facto* standard on graphs, we have attempted to accelerate graph representation learning via importance sampling on large-scale heterogeneous graphs [TKDD20]. The key idea is to design an effective sampler that is aware of the multitude of node and edge types and their complex inter-dependence. Moreover, we have proposed a contrastive graph pre-training approach designed for large-scale heterogeneous graphs [KDD21b], using efficient sampling and sparsification strategies.

III. Future Research Agenda

My future work will still be anchored on learning and mining on graphs. As major research thrusts, I will continue with more in-depth studies in the three areas (areas A, B, C in Figure 1), as many research questions remain open.

For graph representation learning, while we have conducted some preliminary work in trustworthiness and fairness, they are still in their early stages and many challenges exist. First, different perspectives of fairness and different forms of biases exist, and it is crucial to have an adaptive framework that can accommodate different fairness or bias definitions. Second, how do we study various related concepts in trustworthy learning systematically? For example, fairness and explainability are often not independent but have intricate relationships. Considering both or more aspects jointly could improve the trustworthiness of graph learning.

For graph multi-modality, in the age of large models, how can we integrate large amount of visual and/or language data to complement graph learning. Ideally, we could envision a foundation model for graph data, encompassing various modalities. Nevertheless, there would be several fundamental questions regarding the graph foundation model, including its feasibility, its methodology, and its applications.

For data efficiency, I plan to further explore prompt-based learning of graph models beyond current preliminary studies, such as on different kinds of graphs like dynamic graphs, on diverse pre-training models like graph transformers. In addition to prompt-based learning, other parameter-efficient fine-tuning methods are worth exploring on graph models, such as adapter learning and low-rank adaptation techniques.

References in this Statement

A. Graph representation learning

[WSDM11] **Y. Fang** and K. C.-C. Chang. “Searching Patterns for Relation Extraction over the Web: Rediscovering the Pattern-Relation Duality.” In *WSDM* 2011, pp. 825–834.

[SIGIR12] **Y. Fang**, P. Hsu and K. C.-C. Chang. “Confidence-Aware Graph Regularization with Heterogeneous Pairwise Features.” In *SIGIR* 2012, pp. 951–960.

[ICDE13] **Y. Fang**, K. C.-C. Chang and H. W. Lauw. “RoundTripRank: Graph-based Proximity with Importance and Specificity.” In *ICDE* 2013, pp. 613–624.

[ICML14] **Y. Fang,** K. C.-C. Chang and H. W. Lauw. “Graph-based Semi-supervised Learning: Realizing Pointwise Smoothness Probabilistically.” In *ICML* 2014, Part 2, pp. 406–414.

[ICDE16a] **Y. Fang,** V. W. Zheng and K. C.-C. Chang. “Learning to Query: Focused Web Page Harvesting for Entity Aspects.” In *ICDE* 2016, pp. 1002–1013.

[ICDE16b] **Y. Fang**, W. Lin, V. W. Zheng, M. Wu, K. C.-C. Chang and X.-L. Li. “Semantic Proximity Search on Graphs with Metagraph-based Learning.” In *ICDE* 2016, pp. 277–288.

[Methods17] S. Kircali, **Y. Fang**, M. Wu, X. Xiao and X. Li. “Disease Gene Classification with Metagraph Representations.” *Methods* 131:83–92, 2017.

[ICDM18] V. W. Zheng, M. Sha, Y. Li, H. Yang, **Y. Fang**, Z. Zhang, K.-L. Tan and K. C.-C. Chang. “Heterogeneous Embedding Propagation for Large-scale E-Commerce User Alignment.” In *ICDM* 2018, pp. 1434–1439.

[BMC18] S. Kircali, L. Ou-Yang, **Y. Fang**, C.-K. Kwoh, M. Wu and X. Li. “Integrating Node Embeddings and Biological Annotations for Genes to Predict Disease-Gene Associations.” *BMC Systems Biology* 12(Supp 9): 31–44, 2018.

[KDD19] B. Hu, **Y. Fang** and C. Shi. “Adversarial Learning on Heterogeneous Information Networks.” In *KDD* 2019, pp. 120–129.

[TKDE19] **Y. Fang**, W. Lin, V. W. Zheng, M. Wu, J. Shi, K. C.-C. Chang and X. Li. “Metagraph-based Learning on Heterogeneous Graphs.” *IEEE TKDE* 33(1):154–168, 2019.

[TKDE20] W. Zhang, **Y. Fang**, Z. Liu, M. Wu and X. Zhang. “mg2vec: Learning Relationship-Preserving Heterogeneous Graph Representations via Metagraph Embedding.” *IEEE TKDE* 34(3), 2020, pp. 1317–1329.

[SDM20] X. Jiang, B. Hu, **Y. Fang**, C. Shi. “Multiplex Memory Network for Collaborative Filtering.” In *SDM* 2020, pp. 91–99.

[ECMLPKDD20a] Y. Ji, M. Yin, **Y. Fang**, H. Yang, X. Wang, T. Jia and C. Shi. “Temporal Heterogeneous Interaction Graph Embedding For Next-Item Recommendation.” In *ECML-PKDD* 2020, Part III, pp. 314–329.

[IPM20] Y. Ji, C. Shi, **Y. Fang**, X. Kong and M. Yin. Semi-supervised Co-Clustering on Attributed Heterogeneous Information Networks. *IPM* 57(6):102338, 2020.

[SIGIR20] W. Huang, Y. Li, **Y. Fang,** J. Fan and H. Yang. BiANE: Bipartite Attributed Network Embedding.” In *SIGIR* 2020, pp. 149–158.

[TKDD21] S. Ata, **Y. Fang**, M. Wu, J. Shi, C. Kwoh, X. Li. Multi-View Collaborative Network Embedding. In *ACM TKDD* 15(3), 2021.

[JBHI21] Z. Hao, D. Wu, Y. Fang, M. Wu, R. Cai and Xiao-Li Li. Prediction of Synthetic Lethal Interactions in Human Cancers using Multi-view Graph Auto-Encoder. *IEEE Journal of Biomedical & Health Informatics* 25(10), 2021, pp. 4041--4051.

[ECMLPKDD21] Y. Ji, T. Jia, **Y. Fang** and C. Shi. Dynamic Heterogeneous Graph Embedding via Heterogeneous Hawkes Process. In *ECML-PKDD* 2021, Part I, pp. 388–403.

[IJCAI21] Z. Liu, **Y. Fang**, C. Liu and S. C. H. Hoi. Node-wise Localization of Graph Neural Networks. In *IJCAI* 2021, pp. 1520–1526.

[TKDE21a] F. Zhu, **Y. Fang**, K. Zhang, K. C.-C. Chang, H. Cao, Z. Jiang and M. Wu. Unified and Incremental SimRank: Index-free Approximation with Scheduled Principle. In *IEEE TKDE* 35(3), 2021, pp. 3195–3210.

[TKDE21b] Z. Liu, **Y. Fang**, Y. Liu, V. W. Zheng. Neighbor-Anchoring Adversarial Graph Neural Networks. In *IEEE TKDE* 35(1), 2021, pp. 784–795.

[WWW22a] Z. Wen and **Y. Fang**. TREND: TempoRal Event and Node Dynamics for Graph Representation Learning. In *WWW* 2022, pp. 1159–1169.

[IJCAI22] T. Huang, D. Wang, **Y. Fang** and Z. Chen. End-to-End Open-Set Semi-Supervised Node Classification with Out-of-Distribution Detection. In *IJCAI* 2022, pp. 2087–2093.

[WWW23b] T.-K. Nguyen, Z. Liu and **Y. Fang**. Link Prediction on Latent Heterogeneous Graphs. In *WWW* 2023, pp. 263–273.

[AAAI23a] X. Yu, Z. Liu, **Y. Fang** and X. Zhang. Learning to Count Isomorphisms with Graph Neural Networks. In *AAAI* 2023, pp. 4845–4853.

[AAAI23b] Z. Liu, T.-K. Nguyen and **Y. Fang**. On Generalized Degree Fairness in Graph Neural Networks. In *AAAI* 2023, pp. 4525–4533.

[TKDE23] Z. Liu, **Y. Fang** and M. Wu. Dual-View Preference Learning for Adaptive Recommendation. In *IEEE TKDE* 35(11), 2023, pp. 11316–11327.

[TOIS23] Z. Liu, **Y. Fang** and M. Wu. Mitigating Popularity Bias for Users and Items with Fairness-centric Adaptive Recommendation. In *ACM TOIS* 41(3), Article No. 55, 2023.

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

B. Graph multi-modality learning

[TACL14] **Y. Fang** and M.-W. Chang. “Entity Linking on Microblogs with Spatial and Temporal Signals.” *TACL* Vol. 2, October 2014, pp. 259-272. *Invited for oral presentation at EMNLP 2014.*

[IJCAI17] **Y. Fang**, K. Kuan, J. Lin, C. Tan and V. Chandrasekhar. “Object Detection Meets Knowledge Graphs.” In *IJCAI* 2017, pp. 1661–1667.

[CIKM20a] Y.-N. Chuang, C.-M. Chen, C.-J. Wang, M.-F. Tsai, **Y. Fang** and E. P. Lim. TPR: Text-aware Preference Ranking for Recommender Systems. In *CIKM* 2020, pp. 215–224.

[ECMLPKDD20b] Y. Lu, R. Xie, C. Shi, **Y. Fang**, W. Wang, X. Zhang and L. Lin. “Social Influence Attentive Neural Network for Friend-Enhanced Recommendation.” In *ECML-PKDD* 2020, Part IV (Applied Data Science), pp. 3–18.

[CIKM21a] S. Xu, C. Yang, C. Shi, **Y. Fang**, Y. Guo, T. Yang, L. Zhang and M. Hu. Topic-aware Heterogeneous Graph Neural Network for Link Prediction. In *CIKM* 2021, pp. 2261–2270.

[SIGIR23] Z. Wen and **Y. Fang**. Augmenting Low-Resource Text Classification with Graph-Grounded Pre-training and Prompting. In *SIGIR* 2023, pp. 506–516.

C. Data efficiency and scalability on graphs

[VLDB13] F. Zhu, **Y. Fang**, K. C.-C. Chang and J. Ying. "Incremental and Accuracy-Aware Personalized PageRank through Scheduled Approximation." In *VLDB* 2013, pp. 481–492. *Extended version invited to the collection of best papers of VLDB.*

[VLDBJ15] F. Zhu, **Y. Fang**, K. C.-C. Chang and J. Ying. "Scheduled Approximation for Personalized PageRank with Utility-Driven Hub Selection." *VLDBJ* 24(5):655–679, 2015.

[KDD20] Y. Lu, Y. Fang and C. Shi. "Meta-learning on Heterogeneous Information Networks for Cold-start Recommendation." In *KDD* 2020, pp. 1563–1573.

[BIOINF20] R. Cai, X. Chen, **Y. Fang**, M. Wu and Y. Hao. Dual-Dropout Graph Convolutional Network for Predicting Synthetic Lethality in Human Cancers. *Bioinformatics* 36(16):4458–4465, 2020.

[CIKM20b] Z. Liu, W. Zhang, **Y. Fang**, X. Zhang and S. C. H. Hoi. Towards Locality-Aware Meta-Learning of Tail Node Embeddings on Networks. In *CIKM* 2020, pp. 975–984.

[TKDD20] Y. Ji, M. Yin, H. Yang, J. Zhou, V. W. Zheng, C. Shi and **Y. Fang**. Accelerating Large-Scale Heterogeneous Interaction Graph Embedding Learning via Importance Sampling. *ACM TKDD* 15(1), 2020.

[AAAI21a] Y. Lu, X. Jiang, **Y. Fang** and C. Shi. Learning to Pre-train Graph Neural Networks. In *AAAI* 2021, pp. 4276–4284.

[AAAI21b] Z. Liu, **Y. Fang**, C. Liu and S. C. H. Hoi. Relative and Absolute Location Embedding for Few-Shot Node Classification on Graph. In *AAAI* 2021, pp. 4267–4275.

[KDD21a] Z. Liu, K. Nguyen and **Y. Fang**. Tail-GNN: Tail-Node Graph Neural Networks. In *KDD* 2021, pp. 1109–1119.

[KDD21b] X. Jiang, T. Jia, C. Shi, **Y. Fang**, Z. Lin and H. Wang. Pre-training on Large-Scale Heterogeneous Graph. In *KDD* 2021, pp. 756–766.

[SIGIR21] Z. Wen, **Y. Fang** and Z. Liu. Meta-Inductive Node Classification across Graphs. In *SIGIR* 2021, pp. 1219–1228.

[CIKM21b] X. Jiang, Y. Lu, **Y. Fang** and C. Shi. Contrastive Pre-training of GNNs on Heterogeneous Graphs. In *CIKM* 2021, pp. 803–812.

[WWW22b] Z. Liu, Q. Mao, C. Liu, **Y. Fang** and J. Sun. On Size-Oriented Long-Tailed Graph Classification of Graph Neural Networks. In *WWW* 2022, 1506–1516.

[BIOINF22] Y. Long, M. Wu, Y. Liu, **Y. Fang**, C. K. Kwoh, J. Chen, J. Luo and X. Li. Pre-training Graph Neural Networks for Link Prediction in Biomedical Networks. In *Bioinformatics* 38(8), 2022, pp. 2254–2262.

[WWW23a] Z. Liu, X. Yu, **Y. Fang** and X. Zhang. GraphPrompt: Unifying Pre-Training and Downstream Tasks for Graph Neural Networks. In *WWW* 2023, pp. 417–428.

[CIKM23] Z. Wen, **Y. Fang**, Y. Liu, Y. Guo and S. Hao. Voucher Abuse Detection with Prompt-based Fine-tuning on Graph Neural Networks. In *CIKM* 2023 (Applied Research), pp. 4864–4870.

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

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