Neural Shifts in Collaborative Team Recommendation (Half-Day Tutorial)

Mahdis Saeedi
University of Windsor, ON, Canada
msaeedi@uwindsor.ca

Hossein Fani
University of Windsor, ON, Canada
hfani@uwindsor.ca

Abstract

Team recommendation involves selecting skilled experts to form an *almost surely* successful collaborative team, or refining the team composition to maintain or excel at performance. To eschew the tedious and error-prone manual process, various computational and social science theoretical approaches have been proposed wherein the problem definition remains essentially the same, while it has been referred to by such other names as team allocation, selection, composition, and formation. In this tutorial, we study the advancement of computational approaches from greedy search in pioneering works to the recent learning-based approaches, with a particular in-depth exploration of graph neural network-based methods as the cutting-edge class, via unifying definitions, formulations, and evaluation schema. More importantly, we then discuss team *refinement*, a subproblem in team recommendation that involves structural adjustments or expert replacements to enhance team performance in dynamic environments. Finally, we introduce training strategies, benchmarking datasets, and open-source tools, along with future research directions and real-world applications.

```
[10 min] Pioneering Techniques
    [10 min] Subgraph Optimization Objectives and Techniques
[60 min] Paradigm Shift to Neural Team Recommendation
    [10 min] Variational/non-Variational Multilabel Matching
    [25 min] Graph Neural Team Recommendation
    [25 min] End-to-End Approaches
[20 min] Team Refinement
#Break.............................................................
[20 min] Training Strategies
    [05 min] Virtually Unsuccessful Teams via Negative Samplings
    [05 min] Temporal Team Recommendation via Streaming Training
    [10 min] Fair Team Recommendation via Curriculum Learning
[10 min] Evaluation Methodologies
    [05 min] Datasets
    [05 min] Effectiveness/Efficiency
[15 min] Future Directions
    [05 min] Fair & Diverse Team Recommendation
    [05 min] Spatial Team Recommendation
    [05 min] Multi-Objective Neural Team Recommendation
[10 min] Real-World Applications
[35 min] Hands-on OpeNTF & Adila
```

Figure 1: Our tutorial's outline.

1 Motivation

Team recommendation aims to automate forming collaborative teams of experts whose combined skills, applied in coordinated ways, can successfully achieve tasks in real-world scenarios across diverse fields. For instance,

- Medical emergencies need ad-hoc teams of diverse professionals, such as nurses and physicians, to promptly address patient needs, and the challenge lies in automating the selection of the right group of caregivers considering expertise, communication efficiency, availability, and logistical constraints.
- Scientific publication depends on peer reviewers' assessments, yet the process currently remains a bottleneck due to the *manual* process, automating of which based on expertise, past performance, and load balancing, journal editors and conference organizers can reduce administrative overhead, assign more accurate and fair reviewer matches, and accelerate publication cycles.
- Educators need to split students into collaborative teams to enhance peer learning, engagement, and social skill development, which has become of growing interest due to the proliferation of large-scale online classes. By automating team formation among students based on skill diversity, availability, and learning styles, teachers improve group efficacy and equitable collaboration.

Traditionally, teams were composed manually by relying on human experience and instinct, which is a tedious, error-prone, and suboptimal process due to hidden personal and societal biases, a multitude of criteria to optimize, and an overwhelming number of candidates, among other reasons. The team composition can be heavily influenced by subjective opinions that already inherit hidden *un*fair societal biases, largely ignoring the *diversity* in recommended experts and resulting in discrimination and reduced visibility for already disadvantaged experts (e.g., females), disproportionate selection of *popular* experts, and racial disparities. Additionally, since this process involves a multitude of criteria, including project importance, budget, time constraints and team size limitations, the decision-making is all the more difficult and almost impossible for a large-scale pool of experts. Hence, together with business sectors like Linkedin[1], researchers in artificial intelligence and machine learning have long been developing computational models sifting through massive collections of experts and efficiently learning relationships between experts and their skills in the context of successful and *un*successful teams and excel at recommending *almost surely* successful teams.

2 Prior Tutorials

There had been *no* comprehensive tutorial with comparative analysis and critical reviews on the applicability of approaches in real-world scenarios. Recently, we began to fill the gap by providing a tutorial at UMAP24[2] centered on a narrowed scope of subgraph optimization with experts being *online* skilled users, a comprehensive tutorial at SIGIR-AP24[3] based on a novel taxonomy from a computational perspective with a focus on real-world applications, and a technical tutorial at WSDM25[4] to review seminal solutions with a special focus on the emerging graph neural network-based

[1] www.linkedin.com
[2] www.um.org/umap2024/tutorials/
[3] www.sigir-ap.org/sigir-ap-2024/tutorial/
[4] www.wsdm-conference.org/2025/tutorials/

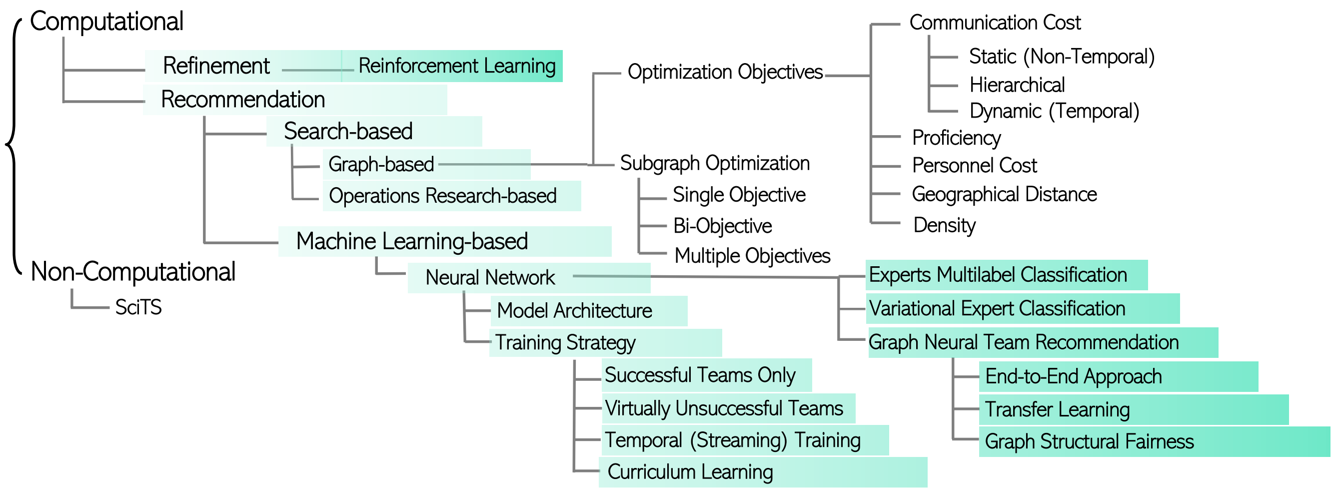

Figure 2: Taxonomy of team recommendation methods.

methods. In this tutorial, from Figure 1 and 2, we expand on them based on our recently published [21] and ongoing surveys of neural approaches that leverage graph neural networks [10, 12, 23], sequence-to-sequence models and transformers [22], and other techniques [23]. Notably, we further investigate team refinement methods, such as expert replacement through reinforcement learning in dynamic sport teams or online video games. We also present proposed methods to mitigate inherent popularity and gender biases via curriculum learning and loss regularization.

3 Target Audience and Knowledge

The team recommendation problem falls under social information retrieval (Social IR), where the *right* group of experts are searched to solve the tasks at hand. In this tutorial, (1) we target *beginner* or *intermediate* researchers, industry technologists and practitioners with a broad interest in developing recommender systems, willing to have a whole picture of team recommendation techniques. (2) Furthermore, we target the graph neural network community for a comprehensive review of subgraph optimization objectives and call them for further development of effective yet efficient graph machine learning with a special focus on team recommendation. Last, having regard to the unified comparative analysis, this tutorial enables (3) organizations and practitioners to compare different models and readily pick the most suitable one for their application to form *almost surely* guaranteed successful teams. The target audience will be those *familiar* with graph theory and neural architectures. Where appropriate, we provide sufficient details about advanced techniques, such as dynamic curriculum learning or graph convolutional networks so that the content will be accessible and understandable to those with a fair knowledge of fundamentals.

4 Proposed Tutorial (`180 minutes`)

We start our tutorial with a brief introduction to the pioneering graph-based team recommendation algorithms based on a taxonomy of computational methods, as shown in Figure 2, then continue to explore the learning-based team recommendation and team refinement methods, focusing on modern methods based on graph neural networks and reinforcement learning.

4.1 Pioneering Techniques (`10 minutes`)

The early computational models for team recommendation were developed in operations research (OR), optimizing objectives using integer linear and/or nonlinear programming (IP). Such work, however, was premised on the mutually independent selection of experts and overlooked the organizational and/or social ties. To bridge the gap, graph-based approaches have been proposed to recommend teams via subgraph optimization where the different aspects of real-world teams are captured like communication cost and geographical proximity [15].

4.2 Neural Shifts (`60 minutes`)

Advances in machine learning, graph neural networks in particular, have led to a paradigm shift toward learning-based methods for team recommendation [20]. These methods are different in that they learn the inherent structure of the ties among experts and their skills. Learning-based methods bring efficiency while enhancing efficacy due to the inherently iterative and online learning procedure, and can address the limitations of subgraph optimization solutions with respect to scalability, as well as dynamic expert networks. In this tutorial, we detail learning-based methods, particularly end-to-end graph neural network-based techniques, which have been building up the following and yielding state-of-the-art performances.

4.2.1 ***Variational/non-Variational Multilabel Classifier.*** Neural team recommendation has started with neural-based multilabel classifiers like a simple feedforward network whose parameters are learned by either maximum likelihood optimization or maximizing a posterior using Bayesian neural models [3, 10].

4.2.2 ***Graph Neural Team Recommendation.*** Next, graph neural networks have received growing attention in the team recommendation problem for their performance on learning the dense vector representation of the skills (Figure 3). The majority of approaches in this category have employed transfer learning techniques that involve pretraining dense vectors independently and feeding them into a neural classifier [3, 12, 20].

4.2.3 ***End-to-End Team Recommendation.*** Following the success of end-to-end graph neural network approaches in tasks such as

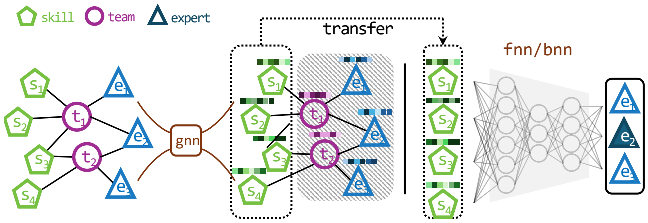

Figure 3: Dense vector representation learning for skills.

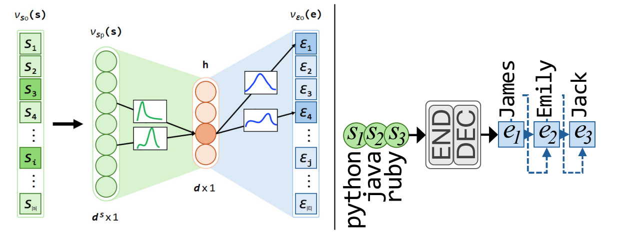

Figure 4: Multilabel vs. seq-to-seq team recommendation.

user-item recommendation and information retrieval, team recommendation studies also proposed end-to-end approaches via graph link predictions between team nodes and expert nodes in an expert collaboration graph. More recently, seq-to-seq and tranformer-based approaches have been proposed, wherein team recommendation has been reformulated as a sequence prediction task that directly maps a set of skills to a set of experts [22] (Figure 4), eschewing the unnecessary step of separately learning the skill node embeddings and transferring them to a neural classifier for team recommendation. Such end-to-end approaches strive to tackle the training challenges of neural classifiers, like the curse of sparsity in the output layer and fragmented learning phases.

4.3 Team Refinement (`20 minutes`)

Reinforcement learning with neural policy estimators has been increasingly employed to learn the dynamics of real-world teams for structural modifications or team member replacements in order to maintain or even improve team effectiveness [4, 13, 24] by modeling the problem in a multi-agent setting where a group of agents synchronize their actions in a decentralized manner within a shared environment to achieve a common goal. The task cannot be completed by any individual agent alone but in a team of agents, and each agent must make a decision regarding which team to join.

4.4 Training Strategies (`20 minutes`)

We explain various strategies to train proposed neural team recommenders, including (1) negative sampling [6] to learn from instances of teams labelled as *virtually* unsuccessful. Benchmark datasets in team recommendation lack unsuccessful teams and researchers proposed different heuristics to draw virtually unsuccessful teams and showed their synergy to the model convergence and improved inference during training and test, respectively; (2) streaming training [8] that encode temporal aspects in team recommendation, and (3) curriculum learning [5] that provide an order between experts from the *easy* popular experts to the *hard* non-popular ones to overcome the neural models' popularity bias (Figure 5).

4.5 Evaluation Strategies (`10 minutes`)

We also discuss the benchmark datasets and what has been considered as successful teams to function as the ground truth. Also, we explore quantitative and qualitative metrics utilized to measure the quality of the recommended teams.

4.6 Challenges and Future Work (`15 minutes`)

We discuss open issues and future directions as follows:

- **Fair and Diverse Teams:** Existing methods primarily focus on maximizing the success rate of recommended teams, largely ignoring *diversity* in the recommended list of experts. Meanwhile, social science research provides compelling evidence about the synergistic effects of diversity on team performance; diversity breeds innovation and increases teams' success by enabling a stronger sense of community and support, reducing conflict, and stimulating more creative thinking. However, there is little work to mitigate biases in team recommender systems [16, 19]. In our tutorial, we introduce notions of fairness and protected attributes and study debiasing algorithms to mitigate the potential unfairness in the models' recommendations.
- **Spatial Team Recommendation:** In search of an optimal team, companies further look for experts in a region where the organization is *geographically* based. Existing methods use skills as a primary factor while overlooking geographical location. We bring forth the *spatial* team recommendation problem; that is, given a set of experts, skills and geolocations, the goal is to determine if the combination of skills and geolocations in forming teams has synergistic effects.
- **Multi-Objective Optimization:** In real-world team recommendation scenarios, balancing multiple, often conflicting objectives (e.g., team effectiveness and experts' workload distribution) requires a training process guided by a loss function that explicitly accounts for multiple objectives. However, existing neural team recommendation approaches, which commonly frame the problem as a classification or link prediction task, aim to maximize the coverage of the required skills and mainly rely on standard loss functions such as cross-entropy, and designing a task-aware loss function is overlooked.

4.7 Applications (`10 minutes`)

We explain novel applications of team recommendation, including:

- **Learning Group Recommendation**: In online classes, where physical presence and interaction are absent, team recommendation connects students to improve their social skills and combat the isolation that can sometimes accompany remote learning [9, 18]. In large classes, where individual interactions with the instructor may be limited, group work ensures that students still have ample opportunities to engage with the material.
- **Reviewer Assignment**: Another application of team recommendation is in peer-review assignments [1, 2] where the reviewers are paired with manuscripts within their expertise for high-quality reviews while managing conflicts of interests [11]. Herein, research topics (skills) and reviewers (experts) are mapped into a latent space, and given a manuscript on a subset of research

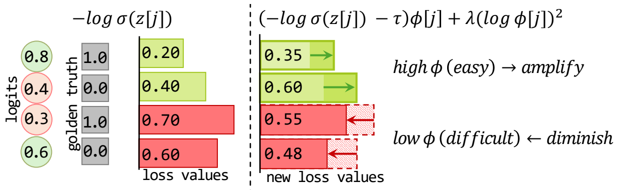

Figure 5: Standard (left) vs. dynamic curriculum (right) loss.

topics, a model recommends top-k optimum reviewers for the research topics.

- **Palliative Care:** Team recommendation is also applicable in healthcare to assign a team of caregivers to patients who seek help with their daily activities due to disease or disorders [14], or to form ad-hoc teams of clinical care experts for medical emergencies [17]. The challenge lies in optimally assigning care providers in teams to address patient needs while considering factors such as communication, distance, and contract costs.

4.8 Hands-On (`35 minutes`)

We introduce publicly available libraries for team recommendation. Notably, we provide hands-on experience with `OpeNTF`[5] [7], an open-source benchmark library for neural models that can efficiently preprocess large-scale datasets, can be easily extended or customized to new neural methods, and is extensible to new datasets from other domains. We also introduce `Adila`[6] [16], which enables further in-processing female-advocate loss regularization [19], as seen in Figure 6, and/or post-processing reranking [16] to the list of recommended experts to ensure a fair outcome. `Adila` is equipped with fairness metrics, which, in tandem with *utility* metrics, allow exploring the trade-offs between fairness and success rate.

5 Presenters (In-Person)[7]

- **Mahdis Saeedi (She/Her)** is a Postdoctoral Fellow at the School of Computer Science and a Lecturer at the Department of Mathematics, University of Windsor. She has taught graph theory, linear algebra, graph representation learning and topics in AI at undergraduate and graduate levels. During her PhD, she focused on theoretical aspects of graph mining, including edge ideals in bipartite and glued graphs. Her current research centers on graph machine learning. She has published in the field's premier venues, such as ACM Computing Surveys, WSDM, ECIR, SIGIR-AP, Transactions of Combinatorics, and the Journal of Algebra.
- **Hossein Fani (He/Him)** is an Assistant Professor at the School of Computer Science, University of Windsor. His research is at the intersection of Social Network Analysis, User Modeling, and Information Retrieval, with a diverse team of `15+` HQP, funded by NSERC-DG, NSERC-RTI, and CFI-JELF. His research appears in ACM Computing Surveys, Elsevier's IP&M, ACM's TOIS, Wiley's JASIST, and SIGIR, CIKM, WSDM and ECIR. He translates his research into techniques for the industry while leading R&D funded by NSERC Alliance and Mitacs Accelerate. He has been granted three patents by the USPTO, including `US10,885,131`, `US11,768,522`, and `US12,067,625`. Fani's teaching experience

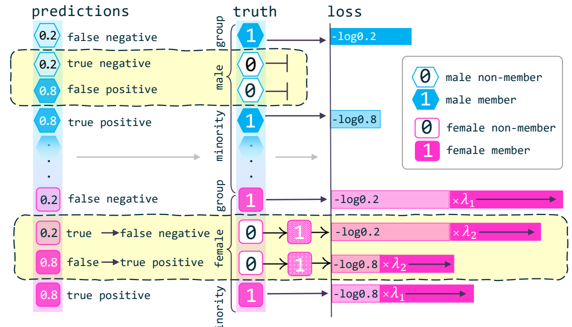

Figure 6: Female-advocate loss regularization.

spans over 15+ years in countries with diverse cultures and educational systems. He has taught courses, including natural language communication and graph representation learning, in multi-section classes ranging from 30 (graduate) to 200+ (undergraduate) with in-person, hyflex, and online modalities. He has developed graduate courses, including big data analytics and the reciprocal role of AI, science, and society.

References

[1] Aksoy et al. 2023. Reviewer Assignment Problem: A Systematic Review of the Literature. *J. Artif. Intell. Res.* 76 (2023), 761–827.

[2] Arabzadeh et al. 2024. Reviewerly: Modeling the Reviewer Assignment Task as an Information Retrieval Problem. In *CIKM 2024*. 5554–5555.

[3] Ahmed et al. 2024. Vector Representation Learning of Skills for Collaborative Team Recommendation: A Comparative Study. In *WISE*.

[4] Bing et al. 2024. Team formation in large organizations: A deep reinforcement learning approach. *Decis. Support Syst.* 187 (2024), 114343.

[5] Barzegar et al. 2025. Adaptive Loss-based Curricula for Neural Team Recommendation. In *WSDM*.

[6] Dashti et al. 2022. Effective Neural Team Formation via Negative Samples. In *CIKM*. ACM, 3908–3912.

[7] Dashti et al. 2022. OpeNTF: A Benchmark Library for Neural Team Formation. In *CIKM*. ACM, 3913–3917.

[8] Fani et al. 2024. A Training Strategy for Future Collaborative Team Prediction. In *ECIR*. Springer.

[9] Gonzalo Candel et al. 2023. An Integer Linear Programming Model for Team Formation in the Classroom with Constraints. In *HAIS*.

[10] Hamidi Rad et al. 2024. A Variational Neural Architecture for Skill-based Team Formation. *ACM Trans. Inf. Syst.* 42, 1 (2024), 7:1–7:28.

[11] Katharina et al. 2023. RevASIDE: Evaluation of Assignments of Suitable Reviewer Sets for Publications from Fixed Candidate Pools. *J. Data Intell.* 4 (2023), 101–133.

[12] Kaw et al. 2023. Transfer Learning with Graph Attention Networks for Team Recommendation. In *IJCNN 2023*. IEEE, 1–8.

[13] Karol Kurach et al. 2020. Google Research Football: A Novel Reinforcement Learning Environment. In *AAAI Conference on Artificial Intelligence*.

[14] Kalyani Selvarajah et al. 2018. Team Formation in Community-Based Palliative Care. In *INISTA*.

[15] Lappas et al. 2009. Finding a team of experts in social networks. In *SIGKDD*. ACM, 467–476.

[16] Loghmani et al. 2023. Bootless Application of Greedy Re-ranking Algorithms in Fair Neural Team Formation. In *ECIR-BIAS 2023*, Vol. 1840. Springer, 108–118.

[17] McLeod et al. 2021. Hacking teamwork in health care: Addressing adverse effects of ad hoc team composition in critical care medicine. *Health care management review* 46, 4 (2021), 341–348.

[18] Michaelsen et al. 2023. *Team-based learning: A transformative use of small groups in college teaching*. Taylor & Francis.

[19] Moasses et al. 2024. vivaFemme: Mitigating Gender Bias in Neural Team Recommendation via Female-Advocate Loss Regularization. In *ECIR-BIAS 2024*.

[20] Nguyen et al. 2023. Learning heterogeneous subgraph representations for team discovery. *Inf. Retr. J.* 26, 1 (2023), 8.

[21] Saeedi et al. 2025. A Survey of Subgraph Optimization for Expert Team Formation. *ACM Comput. Surv.* (May 2025). Just Accepted.

[22] Thang et al. 2025. Translative Neural Team Recommendation: From Multilabel Classification to Sequence Prediction. In *SIGIR*.

[23] Vombatkere et al. 2025. A QUBO Framework for Team Formation. *CoRR* abs/2503.23209 (2025). https://doi.org/10.48550/ARXIV.2503.23209

[24] Yifan Zang et al. 2023. Automatic Grouping for Efficient Cooperative Multi-Agent Reinforcement Learning. In *Advances in Neural Information Processing Systems*.

[5] `https://github.com/fani-lab/OpeNTF`
[6] `https://github.com/fani-lab/Adila`
[7] `fani-lab.github.io/`