DQ4FairIM: Fairness-aware Influence Maximization using Deep Reinforcement Learning

Keywords: Algorithmic Fairness, Fairness-aware Influence Maximization, Reinforcement Learning, Maxmin Fairness, Q-learning

Extended Abstract

The Influence Maximization (IM) is a fundamental problem in social networks that aims to select a set of seed nodes within a given budget to maximize the spread of influence in a social network [1]. However, real-world social networks have several structural inequalities, such as dominant majority groups and underrepresented minority groups, and if these inequalities are not considered while designing IM algorithms, the outcome might be biased, disproportionately benefiting majority groups while marginalizing minorities [2]. To address this challenge, we propose a novel Reinforcement Learning (RL) method, called *DQ4FairIM* (Deep Q-learning for Fair Influence Maximization), that maximizes the expected number of influenced nodes by learning a fair RL policy. The learnt policy ensures that minority groups are not unfairly or disproportionately excluded. The main benefit of using RL for this problem is its generalization ability, wherein a policy learnt on a specific problem instance can be generalized to new problem instances without the need to learn the model from scratch.

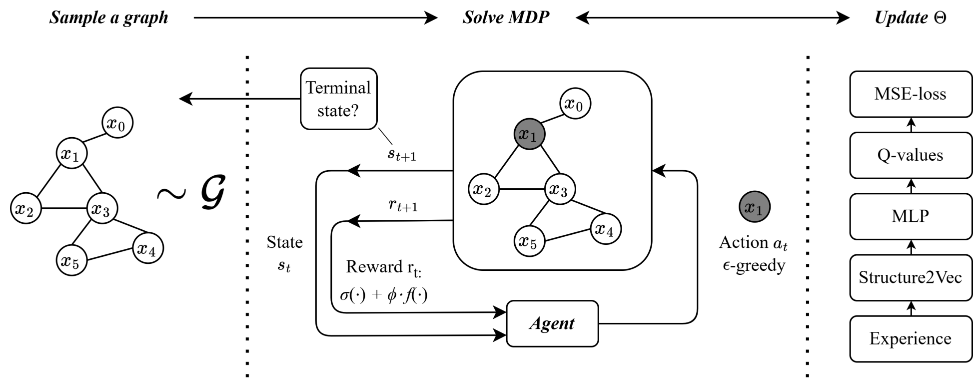

Figure 1: **Overview of DQ4FairIM.** The RL agent iteratively selects seed nodes guided by a reward that balances influence outreach and fairness. Embeddings are computed using Structure2Vec.

We model IM as a Markov Decision Process (MDP), where the RL agent sequentially selects nodes until a budget k is reached (as shown in Figure 1). The reward function combines total outreach $\sigma(G,S)$ with a fairness term $f(G,S)$, defined using the maximin criterion:

$$R(G,S) = \sigma(G,S) + \phi \cdot f(G,S), \tag{1}$$

where $f(G,S)$ is the minimum fraction of influenced nodes across all communities., and ϕ balances the trade-off between outreach and fairness. This ensures that the least-reached group receives higher priority. The agent is trained using Deep Q-learning with Structure2Vec embeddings [3], which capture both graph structure and community membership. An ε-greedy policy ensures balanced exploration and exploitation during training. Importantly, the learned

Table 1: Average outreach and fairness on synthetic and real-world networks.

Method	**Synthetic**						**Real-world**			
	HBA1k		**HBA10k**		**Obesity**		**Facebook**		**Twitter**	
	Out.	Fair.	Out.	Fair.	Out.	Fair.	Out.	Fair.	Out.	Fair.
CELF	0.1857	*0.1832*	0.1043	0.0926	0.0442	0.0356	0.8479	0.8343	*0.1871*	*0.1508*
Degree	0.1815	0.1722	0.1102	0.0976	0.1158	0.1063	0.8414	0.8263	0.1734	0.1294
Pagerank	0.1821	0.1744	0.0975	0.0829	0.1105	0.1011	0.8457	0.8316	0.1742	0.1345
Parity	0.1569	0.1673	0.1097	*0.1007*	0.1151	*0.1113*	0.8413	0.8264	0.1769	0.1364
Fair Pagerank	0.1841	0.1788	*0.1104*	0.0975	0.1103	0.1071	0.8459	0.8317	0.1822	0.1415
CrossWalk	0.1759	0.1697	0.1091	0.1003	0.0899	0.0793	**0.8523**	*0.8381*	0.1822	0.1457
CEA	*0.1870*	0.1817	**0.1119**	0.0999	**0.1182**	0.1075	0.8451	0.8308	**0.1927**	0.1457
DQ4FairIM (ϕ=0)	0.1348	0.1344	0.1062	0.0975	0.1140	0.1027	*0.8520*	0.8378	0.1802	0.1470
DQ4FairIM (ϕ=1)	**0.2098**	**0.2010**	0.1102	**0.1010**	*0.1169*	**0.1146**	**0.8523**	**0.8382**	0.1859	**0.1561**

policy generalizes across unseen networks without retraining which is a key advantage over static methods.

Experiments were conducted on synthetic datasets (Homophily-BA networks, Obesity, and real-world datasets (Facebook and Twitter). Baselines include CELF (Cost-Effective Lazy Forward), Degree, PageRank, Parity, Fair PageRank, CrossWalk[4], and CEA[5]. Results show that DQ4FairIM with $\phi = 1$ consistently achieves the highest fairness across all datasets, while maintaining outreach comparable to or better than existing methods as shown in the Table 1. Under higher propagation probabilities, DQ4FairIM surpasses all baselines in both fairness and outreach. Moreover, the framework generalizes effectively to larger evolving networks, highlighting scalability and robustness in dynamic settings.

Ethical Considerations: Our work is based entirely on publicly available datasets and does not involve the collection of personal or sensitive information. The central ethical consideration is ensuring fairness in influence maximization, so that information spread does not disproportionately favor majority groups while marginalizing minorities. The approach is designed to promote equitable information access and reduce disparities in social capital within social networks.

References

[1] Suman Banerjee, Mamata Jenamani, and Dilip Kumar Pratihar. "A survey on influence maximization in a social network". In: *Knowledge and Information Systems* 62.9 (2020), pp. 3417–3455.

[2] Akrati Saxena, George Fletcher, and Mykola Pechenizkiy. "Fairsna: Algorithmic fairness in social network analysis". In: *ACM Computing Surveys* 56.8 (2024), pp. 1–45.

[3] Elias Khalil et al. "Learning combinatorial optimization algorithms over graphs". In: *Advances in neural information processing systems* 30 (2017).

[4] Ahmad Khajehnejad et al. "Crosswalk: Fairness-enhanced node representation learning". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36. 11. 2022, pp. 11963–11970.

[5] Kaicong Ma et al. "Fair influence maximization in social networks: A community-based evolutionary algorithm". In: *IEEE Transactions on Emerging Topics in Computing* (2024).