Adaptive Multi-Round Influence Maximization with Limited Information

Extended Abstract

Vincenzo Auletta
Università degli Studi di Salerno
Fisciano (SA), Italy
auletta@unisa.it

Francesco Carbone
Università degli Studi di Salerno
Fisciano (SA), Italy
frcarbone@unisa.it

Diodato Ferraioli
Università degli Studi di Salerno
Fisciano (SA), Italy
dferraioli@unisa.it

Cosimo Vinci
Università del Salento
Lecce, Italy
cosimo.vinci@unisalento.it

ABSTRACT

The Influence Maximization problem is a classic and well-studied problem in the area of Social Networks Analysis. In this problem you have a social network, a given information diffusion model, and a budget B, and you have to select a set of at most B nodes (seeds) to activate in order to start an information diffusion campaign that is able to reach the (expected) largest number of nodes in the network. Recently, to better model viral marketing scenarios where advertisers conduct multiple rounds of viral marketing to promote one product, attention has been given to the adaptive and the multi-round versions of the problem. Here the campaign is orchestrated on a horizon of T rounds and at the beginning of each round a different set of seeds is activated that can be adaptively selected given the results of the previous rounds. In this paper we generalize this setting to the case where the diffusion probabilities of the links in the network are not known in advance and they have to be learned while the campaign is running.

We study the problem under the lens of online bandit algorithms, and we propose an online learning algorithm that is able to achieve a constant approximation of the optimal solution with only constant regret with respect to T. We also propose an alternative approach and we give preliminary experimental evidence that this outperforms our online learning algorithm in terms of computational complexity, keeping the regret sublinear.

CCS CONCEPTS

• **Theory of computation → Online learning algorithms**; **Social networks**.

KEYWORDS

Influence Maximization; Online Learning; Adaptive Algorithms; Greedy Algorithms

This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

ACM Reference Format:
Vincenzo Auletta, Francesco Carbone, Diodato Ferraioli, and Cosimo Vinci. 2025. Adaptive Multi-Round Influence Maximization with Limited Information: Extended Abstract. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

EXTENDED ABSTRACT

In the last decades online social networks are getting more and more popular as a channel for sharing and searching information. As a consequence, these networks are now the favorite channel for marketing or political campaigns, with the aim of influencing people's opinions and choices toward certain goals. The crucial problem that influencers, advertisers and social media managers have to deal with in designing their social campaigns is how to maximize their influence taking advantage of the information diffusion properties of their social networks. A popular marketing technique is to select a (limited) number of starting nodes (seeds) from which to start an information diffusion campaign taking advantage of the word-of-mouth phenomenon. Thus, designers are interested in algorithms that select their seeds in order to maximize the (expected) number of nodes that will be reached by the campaign. The problem turns out to have several applications in marketing [15, 21, 22], opinion formation [3–5, 18], voting [6, 9, 28], and health prevention [29].

In the seminal paper by Kempe et al [22] the *Influence Maximization* problem has been formalized as follows: a social network given in input represents the social relationships among the agents together with the strength of their relations (expressed as probability of success of the information diffusion among these nodes), and a set of at most B seeds has to be selected from which to start an information diffusion campaign. In the following years a plethora of follow-up models have been proposed to extend this basic model along several different directions. In particular, in this work we focus on three extensions that have received great attention: multi-round campaigns, adaptivity, and partial knowledge of the network.

Adaptivity is a powerful technique in the optimization of stochastic problems that recently has received large attention. In the setting of information maximization problems, adaptive algorithms can select seeds sequentially and the i–th seed is selected only after having observed all the nodes reached by the previously chosen seeds. The analysis of optimization advantages of adaptive w.r.t. non-adaptive algorithms has been initiated by Dean et al. [13, 14] on classical packing problems. Only recently, this kind of analysis has been applied to generalizations of the influence maximization problem [7, 10, 11, 20, 24].

Research in influence maximization focused almost exclusively on the *single-round* setting. However, it has been observed that in several real-world examples the influence maximization process works in several rounds even if the budget on the number of selectable seeds is allocated for the entire campaign. This has been observed in electoral campaigns [16, 23], in hiring campaigns [25], or in (viral) advertising where advertisers provide a total budget to brokers, that will allocate it over multiple rounds through suitable budget pacing algorithms [17].

Auletta et al. [2] proposed the *Multi-round Adaptive Information Maximization* problem in which the budget on the number of seeds is defined over multiple rounds but there is no limit on the fraction of budget used in each round, and in each round seeds are chosen adaptively. The goal is then to allocate the seeds along the different rounds in order to maximize the (expected) sum over all rounds of nodes that are infected. Auletta et al. [2] follow the classical approach of Kempe et al. [22], and it assumes that the campaign designer knows the whole network and, in particular, they exactly know the probability that an agent is successful in disseminating information over her social relations. This is a quite unrealistic assumption and in several real-world scenarios the designer has to design the campaign by having a limited information about the structure of the network.

To overcome this limit the Influence Maximization problem has been studied in the framework of online learning algorithms, aka bandit algorithms [8, 26, 27]. Here, the designers have to select the seeds and run the influence maximization campaign while they are learning the diffusion probabilities. As usual in the online learning setting, our algorithm is evaluated with respect to both the approximation of the optimal solution that this algorithm is able to guarantee, and on the speed with which this algorithm approaches to this goal, as measured by the *regret*, that is the difference between the objective solution as reached by an optimization algorithm and the solution reached by the online learning algorithm.

We remark that all these three lines of research have been (separately) explored in state-of.the-art literature. However, to the best of our knowledge, they were never considered collectively. Here, instead, we design online learning algorithms for the *Multi-round Adaptive Influence Maximization with Limited Information* problem.

Our Contribution. In this work we provide a polynomial-time algorithm that is able to achieve, for every $\varepsilon > 0$, a $[\frac{1}{2}(1 - \frac{1}{e}) - \varepsilon]$-approximation of the total number of nodes influenced by the optimal fully-informed adaptive algorithm with constant regret with respect to the time horizon. Note that the approximation ratio matches the one of the best known polynomial approximation fully-informed algorithm for the problem [2]. As for the regret, a constant dependence on the time horizon is the best that can be achieved. Moreover, this largely improves upon the poly-logarithmic regret achieved in previous works for restricted or similar settings [12, 19]. Thus, we are able to exploit the absence of per-round budget limits to learn very quickly the underlying probability distributions.

Interestingly, this result is built on a technical result that may be of independent interest. Indeed, Auletta et al. [2] presented an algorithm that, applied to our setting with a budget B and a time horizon T, adaptively and greedily chooses approximately B/T seeds at each round, with each selected seed being with high probability a δ-additive approximation of the seed maximizing the expected increment on the objective function according to the known real diffusion probabilities among nodes. This algorithm has been proved to be a $\frac{1}{2}\left(1 - \frac{1}{e} - \varepsilon\right)$-approximation of the optimal algorithm in the fully-informed setting. We here observe that a good approximation of the optimal seed can be computed even if we do not know real probabilities, but we have only a close approximation to them. This then allows to frame the above algorithm in a simple online learning framework as follows: do exploration as long as we have a close approximation of real diffusion probabilities, and then exploits them as described above. We will prove that this framework allows to essentially match the approximation guarantee of the fully-information algorithm, with a regret that depends only on the length of the exploration phase and the fraction of budget consumed during this phase.

In order to validate the utility of our framework, we apply it to evaluate two alternative policies. For both of them, we show that we can learn diffusion probabilities in a number of steps and with a consumption of budget that does not depend on the time horizon T, achieving in this way a constant regret. The two policies differ on the effective value of the regret, and on the amount of seeds chosen in the exploration phase: indeed, the first policy achieves a better regret but it can require to choose many agents (even all of them) as seeds in each round of the exploration phase; the second policy requires a more moderate seeding in the exploration phase at the cost of a slight increment in the value of the regret. We stress that further applications of our framework can still be developed.

As suggested above, our policies need to compute a sufficiently good estimation of the expected increment in the objective function guaranteed by each possible seed choice. We observe that this can be done through a polynomial number of Monte Carlo simulations, and hence our policies have polynomial time complexity. However, it is well known that Monte Carlo simulations can be very expensive, and this can reduce the practical adoption of our policies when we need to run them over large networks.

To address this issue, we here present an alternative faster policy, based on the UCB framework [1, 19], that keeps upper confidence bounds on the value of n^2 variables, namely the probabilities that information starting from a seed u will infect a node v, (and thus it does not need to estimate them). This policy guarantees a regret that is sublinear with respect to the time horizon T, namely $O(T^{2/3}\sqrt[3]{\log T})$, when $B \leq T$. Moreover, we provide preliminary experimental evidence that this policy still provides similar guarantees even in the case that $B > T$.

ACKNOWLEDGMENTS

This work is partially supported by: the PNRR MIUR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI; MUR - PNRR IF Agro@intesa; the Project SERICS (PE00000014) under the NRRP MUR program funded by the EU – NGEU; GNCS-INdAM.

REFERENCES

[1] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Mach. Learn.* 47, 2-3 (2002), 235–256.
[2] Vincenzo Auletta, Diodato Ferraioli, and Cosimo Vinci. 2024. Stochastic Multi-round Submodular Optimization with Budget. arXiv:2404.13737 [cs.DS]
[3] Vittorio Bilò, Diodato Ferraioli, and Cosimo Vinci. 2022. General Opinion Formation Games with Social Group Membership. In *IJCAI.* 88–94.
[4] David Bindel, Jon M. Kleinberg, and Sigal Oren. 2015. How bad is forming your own opinion? *Games Econ. Behav.* 92 (2015), 248–265.
[5] Robert Bredereck, Lilian Jacobs, and Leon Kellerhals. 2020. Maximizing the Spread of an Opinion in Few Steps: Opinion Diffusion in Non-Binary Networks. In *IJCAI.* 1622–1628.
[6] Matteo Castiglioni, Diodato Ferraioli, Nicola Gatti, and Giulia Landriani. 2021. Election manipulation on social networks: Seeding, edge removal, edge addition. *Journal of Artificial Intelligence Research* 71 (2021), 1049–1090.
[7] Wei Chen and Binghui Peng. 2019. On Adaptivity Gaps of Influence Maximization Under the Independent Cascade Model with Full-Adoption Feedback. In *ISAAC.* 24:1–24:19.
[8] Wei Chen, Yajun Wang, Yang Yuan, and Qinshi Wang. 2016. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *The Journal of Machine Learning Research* 17, 1 (2016), 1746–1778.
[9] Federico Corò, Emilio Cruciani, Gianlorenzo D'Angelo, and Stefano Ponziani. 2019. Exploiting Social Influence to Control Elections Based on Scoring Rules. In *IJCAI.* 201–207.
[10] Gianlorenzo D'Angelo, Debashmita Poddar, and Cosimo Vinci. 2021. Improved Approximation Factor for Adaptive Influence Maximization via Simple Greedy Strategies. In *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021 (LIPIcs, Vol. 198).* 59:1–59:19.
[11] Gianlorenzo D'Angelo, Debashmita Poddar, and Cosimo Vinci. 2023. Better bounds on the adaptivity gap of influence maximization under full-adoption feedback. *Artif. Intell.* 318 (2023), 103895.
[12] Debojit Das, Shweta Jain, and Sujit Gujar. 2022. Budgeted Combinatorial Multi-Armed Bandits. In *AAMAS.* 345–353.
[13] Brian C Dean, Michel X Goemans, and Jan Vondrák. 2005. Adaptivity and approximation for stochastic packing problems. In *SODA.* 395–404.
[14] Brian C Dean, Michel X Goemans, and Jan Vondrák. 2008. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Mathematics of Operations Research* 33, 4 (2008), 945–964.
[15] Pedro M. Domingos and Matthew Richardson. 2001. Mining the network value of customers. In *SIGKDD.* ACM, 57–66.
[16] Elin Falguera, Samuel Jones, and Magnus Ohman. 2014. *Funding of Political Parties and Election Campaigns: A Handbook on Political Finance.* 458 pages.
[17] Joaquin Fernandez-Tapia. 2015. Optimal budget-pacing for Real-time Bidding. *Social Science Electronic Publishing* (2015).
[18] Diodato Ferraioli, Paul W. Goldberg, and Carmine Ventre. 2016. Decentralized dynamics for finite opinion games. *Theor. Comput. Sci.* 648 (2016), 96–115.
[19] Victor Gabillon, Branislav Kveton, Zheng Wen, Brian Eriksson, and Shanmugavelayutham Muthukrishnan. 2013. Adaptive submodular maximization in bandit setting. *NeurIPS* 26 (2013).
[20] Daniel Golovin and Andreas Krause. 2011. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research* 42 (2011), 427–486.
[21] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the spread of influence through a social network. In *KDD.* 137–146.
[22] David Kempe, Jon M. Kleinberg, and Éva Tardos. 2005. Influential Nodes in a Diffusion Model for Social Networks. In *ICALP.* 1127–1138.
[23] OECD. 2016. *Financing Democracy.* 208 pages.
[24] Binghui Peng and Wei Chen. 2019. Adaptive influence maximization with myopic feedback. *NeurIPS* 32 (2019).
[25] E Schmitz, Michael J Moskowitz, David Gregory, and David Reese. 2008. Recruiting Budgets, Recruit Quality, and Enlisted Performance. *CNA Research Memorandum D* 1703 (2008), 5.
[26] Sharan Vaswani, Laks Lakshmanan, Mark Schmidt, et al. 2015. Influence maximization with bandits. *arXiv preprint arXiv:1503.00024* (2015).
[27] Zheng Wen, Branislav Kveton, Michal Valko, and Sharan Vaswani. 2017. Online influence maximization under independent cascade model with semi-bandit feedback. *NeurIPS* 30 (2017).
[28] Bryan Wilder and Yevgeniy Vorobeychik. 2018. Controlling Elections through Social Influence. In *AAMAS.* 265–273.
[29] Amulya Yadav, Bryan Wilder, Eric Rice, Robin Petering, Jaih Craddock, Amanda Yoshioka-Maxwell, Mary Hemler, Laura Onasch-Vera, Milind Tambe, and Darlene Woo. 2017. Influence maximization in the field: The arduous journey from emerging to deployed application. In *AAMAS.* 150–158.