Thompson Sampling with Cumulative Oversampling for Budgeted Influence Maximization

Zhen Xu[1], Xinjie Xing[1], and Shatian Wang[2]

[1]Management School, University of Liverpool, Liverpool, United Kingdom,
[2]IEOR, Columbia University, New York City, United States,

This work studies online budgeted influence maximization and establishes new results on both the offline benchmark and online learning algorithm design for this class of NP-hard problems. The contributions consist of a near-optimal offline approximation framework and a Thompson Sampling–based online algorithm with tight regret guarantees.

As an offline benchmark, the paper introduces *Randomized Budgeted Influence Maximization (R-IMB)*, which allows randomized seed-set selection under an expected budget constraint. The optimal value of R-IMB upper-bounds the achievable average reward of any online policy over a finite horizon. Based on this benchmark, the paper proposes an approximation oracle that achieves a $(1 - 1/e)$ approximation ratio, matching the optimal guarantee for classical influence maximization [5], [6] and improving upon the best-known guarantees for deterministic budgeted influence maximization, which are bounded by $1 - 1/\sqrt{e}$ [7]. The oracle supports heterogeneous node costs and is compatible with scalable influence spread estimation using Reverse Reachable (RR) sets [1], [8]. A concave error interval analysis guarantees that, with high probability, influence spread is uniformly approximated over all feasible seed sets, yielding a $(1 - 1/e - \varepsilon)$ approximation with near-linear runtime complexity.

For the online setting, the paper develops *TS-CO*, a Thompson Sampling algorithm based on a novel *cumulative oversampling* mechanism. Prior work on online influence maximization with learning has largely relied on UCB-style algorithms [3], [9], which admit strong regret bounds but often underperform empirically. Although Thompson Sampling has demonstrated superior empirical performance in many bandit settings [2], [4], its theoretical analysis for combinatorial problems with NP-hard offline structure has remained limited.

Cumulative oversampling constructs optimistic estimates of influence probabilities by aggregating posterior samples across rounds. In each round, a single Gaussian sample is drawn around the regularized least-squares estimator, and cumulative samples are combined to form an estimator that concentrates at logarithmic rates and becomes an upper confidence bound with probability converging to one. When coupled with the R-IMB approximation oracle, TS-CO achieves a scaled regret bound of order $\tilde{O}(dC\sqrt{mT})$, where d is the feature dimension, m the number of edges, and C a network-dependent complexity parameter. This matches the best-known regret bounds for UCB-based algorithms with linear generalization [9] and constitutes the first regret guarantee for a Thompson Sampling–based algorithm for influence maximization that avoids linear dependence on the inverse minimum edge observation probability.

References

[1] Christian Borgs, Michael Brautbar, Jennifer T. Chayes, and Brendan Lucier. Influence maximization in social networks: Towards an optimal algorithmic solution. *CoRR*, abs/1212.0884, 2012.

[2] Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 2249–2257. Curran Associates, Inc., 2011.

[3] Wei Chen, Yajun Wang, and Yang Yuan. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *CoRR*, abs/1407.8339, 2014.

[4] Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In Nader H. Bshouty, Gilles Stoltz, Nicolas Vayatis, and Thomas Zeugmann, editors, *Algorithmic Learning Theory*, pages 199–213, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[5] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '03, pages 137–146, New York, NY, USA, 2003. ACM.

[6] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*, 14(1):265–294, Dec 1978.

[7] Huy Nguyen and Rong Zheng. Influence spread in large-scale social networks - A belief propagation approach. *CoRR*, abs/1204.4491, 2012.

[8] Youze Tang, Xiaokui Xiao, and Yanchen Shi. Influence maximization: Near-optimal time complexity meets practical efficiency. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, pages 75–86, New York, NY, USA, 2014. ACM.

[9] Zheng Wen, Branislav Kveton, Michal Valko, and Sharan Vaswani. Online influence maximization under independent cascade model with semi-bandit feedback. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 3022–3032. Curran Associates, Inc., 2017.