# On Augmented Stochastic Submodular Optimization: Adaptivity, Multi-Rounds, Budgeted, and Robustness

Vincenzo Auletta<sup>1,†</sup>, Diodato Ferraioli<sup>1,\*,†</sup> and Cosimo Vinci<sup>1,2,†</sup>

<sup>1</sup>Università degli Studi di Salerno, Fisciano, Salerno, Italy <sup>2</sup>Università del Salento, Lecce, Italy

#### Abstract

In this work we consider the problem of Stochastic Submodular Maximization, in which we would like to maximize the value of a monotone and submodular objective function, subject to the fact that the values of this function depend on the realization of stochastic events. This problem has applications in several areas, and in particular it well models basic problems such as influence maximization and stochastic probing.

In this work, we advocate the necessity to extend the study of this problem in order to include several different features such as a budget constraint on the number of observations, the chance of adaptively choosing what we observe or the presence of multiple rounds. We here speculate on the possible directions that this line of research can take. In particular, we will discuss about interesting open problems mainly in the settings of robust optimization and online learning.

#### Keywords

Influence Maximization, Stochastic Probing, Approximation Algorithms, Online Learning

## Introduction

Consider the well-known problem of influence maximization: here we are given a (social) network, with edge probabilities encoding the strength of the relationships among network's components; these probabilities describe how easily influence (e.g., an information or an innovation) spreads over the network starting from a limited set of nodes, usually termed *seeds*; the goal is then to find the set of seeds that maximizes the number of "influenced" nodes. The non-adaptive version of the problem, in which seeds are chosen before that influence starts, has been widely studied, with numerous successful algorithms proposed [1, 2, 3] and applications ranging from viral marketing [1], to election manipulation [4, 5], to infective disease prevention [6]. Recently, focus moved to the adaptive variant of this problem (in which seeds can be added even when the influence is spreading from previously placed seeds) [7, 8, 9, 10].

IPS-RiCeRcA-SPIRIT 2022: 10th Italian Workshop on Planning and Scheduling, RiCeRcA Italian Workshop, and SPIRIT Workshop on Strategies, Prediction, Interaction, and Reasoning in Italy.

<sup>\*</sup>Corresponding author.

<sup>&</sup>lt;sup>†</sup>These authors contributed equally.

 <sup>▲</sup> auletta@unisa.it (V. Auletta); dferraioli@unisa.it (D. Ferraioli); cosimo.vinci@unisalento.it (C. Vinci)
♦ https://docenti.unisa.it/001366/home (V. Auletta); https://docenti.unisa.it/023604/home (D. Ferraioli); https://www.unisalento.it/scheda-utente/-/people/cosimo.vinci (C. Vinci)

 <sup>0000-0002-7875-3366 (</sup>V. Auletta); 0000-0002-7962-5200 (D. Ferraioli); 0000-0001-7741-9342 (C. Vinci)
2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

This variant has strong relationship with some *Stochastic Probing* problems [11], in which there are n items whose states are stochastically defined according to (not necessarily equal) independent distributions and the goal is to choose a feasible subset of items maximizing the value of a non-negative, monotone and submodular set function. Several stochastic probing settings can be found in expert hiring [12], online dating and kidney exchange [13], Bayesian mechanism design [14], and sequentially posted pricing [15, 16]. Research both in influence maximization and stochastic probing mainly focused on the single-round setting. However, in several cases the optimization is run in a multi-round setting. For example, viral marketing or electoral campaigns are usually deployed in multiple rounds of (usually adaptive) influence; similarly, expert hiring sessions and sequentially posted pricing are run in multiple steps.

In this paper, we then advocate the necessity to extend the study of these problems in order to consider all these properties: budget, adaptivity, and multiple-rounds. Indeed, previous works only consider them separately.

Indeed, the analysis of optimization advantages of adaptive w.r.t. non-adaptive algorithms has been recently applied to stochastic maximization problems, including *single-round* influence maximization and stochastic probing. In particular, Golovin and Krause [7] showed that whenever the objective function satisfies a property, named *adaptive submodularity*, then a greedy algorithm achieves a constant approximation of the optimal greedy algorithm. Some results have been recently presented about (Adaptive) Multi-round Influence Maximization [17, 18]: these results however do not take into account budget.

Multiple rounds have been considered mainly in the setting of *online learning algorithms*, aka bandit algorithms, in which our algorithm are asked to work even if the input is only partially known, and to improve their performance over multiple round (by essentially learning the input). This class of algorithms has been introduced by Gittins [19], and it has been applied recently also to influence maximization and other submodular maximization processes [20, 21, 22, 23, 24]. While these works focus on multi-round settings, they do not consider the chance of adaptively choose observations, and often they also do not consider budget constraints.

## **Our contribution**

In this work, we introduce the problem of *Stochastic Budgeted Multi-round Adaptive Submodular Maximization*. The problem is defined by a number T of rounds, a ground set of n items, whose state is drawn at each time step from a (node-specific) set  $\Theta$  according to some time-dependent distribution, a time-dependent monotone and submodular state function that assigns a nonnegative real-value to each realization of states upon specific item observations, and a budget B on the possible observations done within the time-horizon T. The problem is then to choose which observations to make in each of the T rounds, subject to the budget limit, in order to maximize the expected total value of the state function over the T rounds. It is not hard to see how this problem generalizes both the influence maximization and stochastic probing problems: in the first case, the state of an item v defines which nodes would be influenced whether v is selected as a seed, with the distribution of the state of v depending on edges' probabilities; an observation consists on choosing a given node as a seed, then running the influence spread, and finally evaluating the realized state by, for example, counting the number or the total weight of influenced nodes. In the second case, the state of an item can be active or inactive, and we can observe the state of selected items only; then, the value of the submodular set function depends on the set of selected items which are active. Note that SBMSm not only generalizes influence maximization and stochastic probing problems by extending them to multiple round settings, but it also allows to model more realistic settings where states can change over time: e.g., in the influence maximization setting, we may allow edge probabilities or node weights to change over time.

Nevertheless the complexity of the problem, we claim that the state of our knowledge is sufficient to face this kind of problems. In fact, we can prove that, by greedily allocating each unit of budget to the round in which it is expected to produce the largest expected value of the score function, given the amount of budget still available, and then running one of the well-known single-round adaptive budgeted approaches to choose observations to make at each round, we may achieve a constant approximation of the optimal algorithm (that adaptively allocates the budget among rounds, and the allocations at each round).

Let us now discuss few directions for the research about Stochastic Budgeted Multi-round Adaptive Submodular Maximization that we deem as very interesting and promising. In particular, we will henceforth focus on two main approaches.

**Online Learning.** Suppose that state distributions are unknown but all rounds share the same state distribution. Can we define an adaptive policy that maximizes the expected total value of the state function by "learning" the underlying probability distribution? In particular, we will consider an *agnostic (adaptive multi-round) policy*, where the choice of each picked item v is based on the observed partial state (both in the current round and the previous ones), on the initial instance, but not on the distribution of states. As each observation provides samples of several local states, such samples will be used to determine an *estimated distribution* of states at each round t, that is more accurate as t increases and more samples are provided. The estimated distribution will be used to "simulate" the real one with a good approximation (e.g., when computing the item maximizing the expected increment).

As benchmark quality metric for the proposed policy  $\pi$ , we will consider the  $\alpha$ -approximate regret  $Reg_{\alpha}(T) = \alpha \cdot \sigma_{\mu}(\pi^*) - \sigma_{\mu}(\pi)$ , where  $\pi^*$  is an optimal non-agnostic adaptive policy (based on the knowledge of the distribution of states), and  $\sigma_{\mu}(\pi)$  denotes the expected value of  $\pi$ .

We believe it would be interesting to evaluate whether there is a policy  $\pi$  with constant approximation and sublinear regret. In particular, it has been recently introduced an approach, named *Combinatorial Upper Confidence Bound* (CUCB), to design online learning algorithm with constant approximation and poly-logarithmic regret [20, 21, 22], that has been successfully used by Gabillon et al. [23] and by Das et al. [24] in setting that are similar but simpler than the one proposed in this work. Unfortunately, it is not hard to see that a naive application of CUCB to the model defined above cannot work, since we need not only to estimate which observations to make at each round, but also how many, and the latter can be largely affected by an imprecise estimation of states' distribution. However, this raises other interesting questions: are there settings in which we can be able to use the CUCB framework to achieve constant approximation and poly-logarithmic regret? Can we design another framework such that is able to return constant approximation and poly-logarithmic regret in our more general setting? Or, how high is the cost that we have to pay in term of approximation guarantee in order to achieve poly-logarithmic regret?

While the regret is an useful measure (and the standard de facto) in order to understand the learning abilities of an algorithm, other measures can be of some interest. For example, one may be interested in agnostic policies that are able to reach a given goal (e.g., the aggregate state function is above a given threshold) within a given deadline. Note that learning algorithm with low regret may fail in achieving this goal: indeed, regret may depend on other time-independent factors that may be often large. Moreover, low regret does not exclude long *exploration* phases in which a lot of "value" has been lost.

**Robust Optimization.** The second approach is also motivated by partial or imprecise knowledge of state distributions. However, while above we aimed to learn these distributions, we here are concerned on how and how much this limited knowledge affects our optimization goals. This approach is receiving large attention very recently, motivated by the necessity to run optimization algorithms not on real data but on predictions often generated by machine learning algorithms. The field has blossomed with applications to facility location [25], scheduling [26], clustering [27], influence maximization [28].

Here, natural interesting questions that we can ask are: how much the known optimization frameworks are sensible to distortion on input data? Or how much has to be lost in that framework if one want an algorithm that is robust to distortions within a given threshold?

### Acknowledgments

This work is partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Markets", by "GNCS–INdAM" and by the PON R&I 2014-2020 Project TEBAKA "Sistema per acquisizione conoscenze di base del territorio".

# References

- [1] D. Kempe, J. Kleinberg, É. Tardos, Maximizing the spread of influence through a social network, in: KDD, 2003, pp. 137–146.
- [2] N. Chen, On the approximability of influence in social networks, SIAM Journal on Discrete Mathematics 23 (2009) 1400–1415.
- [3] C. Borgs, M. Brautbar, J. Chayes, B. Lucier, Maximizing social influence in nearly optimal time, in: SODA, 2014, pp. 946–957.
- [4] B. Wilder, Y. Vorobeychik, Controlling elections through social influence, in: AAMAS, 2018, pp. 265–273.
- [5] M. Castiglioni, D. Ferraioli, N. Gatti, G. Landriani, Election manipulation on social networks: Seeding, edge removal, edge addition, Journal of Artificial Intelligence Research 71 (2021) 1049–1090.
- [6] A. Yadav, B. Wilder, E. Rice, R. Petering, J. Craddock, A. Yoshioka-Maxwell, M. Hemler,

L. Onasch-Vera, M. Tambe, D. Woo, Influence maximization in the field: The arduous journey from emerging to deployed application, in: AAMAS, 2017, pp. 150–158.

- [7] D. Golovin, A. Krause, Adaptive submodularity: Theory and applications in active learning and stochastic optimization, Journal of Artificial Intelligence Research 42 (2011) 427–486.
- [8] W. Chen, B. Peng, On adaptivity gaps of influence maximization under the independent cascade model with full-adoption feedback, in: ISAAC, 2019, pp. 24:1–24:19.
- [9] G. D'Angelo, D. Poddar, C. Vinci, Better bounds on the adaptivity gap of influence maximization under full-adoption feedback, in: AAAI, 2021, pp. 12069–12077.
- [10] G. D'Angelo, D. Poddar, C. Vinci, Improved approximation factor for adaptive influence maximization via simple greedy strategies, in: ICALP, 2021, pp. 59:1–59:19.
- [11] A. Asadpour, H. Nazerzadeh, A. Saberi, Stochastic submodular maximization, in: WINE, 2008, pp. 477–489.
- [12] M. Hoefer, K. Schewior, D. Schmand, Stochastic probing with increasing precision, in: IJCAI, 2021, pp. 4069–4075.
- [13] A. Gupta, V. Nagarajan, A stochastic probing problem with applications, in: IPCO, 2013, pp. 205–216.
- [14] A. Asadpour, H. Nazerzadeh, Maximizing stochastic monotone submodular functions, Management Science 62 (2016) 2374–2391.
- [15] S. Chawla, J. D. Hartline, D. L. Malec, B. Sivan, Multi-parameter mechanism design and sequential posted pricing, in: STOC, 2010, pp. 311–320.
- [16] M. Adamczyk, A. Borodin, D. Ferraioli, B. d. Keijzer, S. Leonardi, Sequential posted price mechanisms with correlated valuations, in: WINE, 2015, pp. 1–15.
- [17] L. Sun, W. Huang, P. S. Yu, W. Chen, Multi-round influence maximization, in: KDD, 2018, pp. 2249–2258.
- [18] G. Tong, R. Wang, Z. Dong, X. Li, Time-constrained adaptive influence maximization, IEEE Transactions on Computational Social Systems 8 (2020) 33–44.
- [19] J. C. Gittins, Bandit processes and dynamic allocation indices, Journal of the Royal Statistical Society: Series B (Methodological) 41 (1979) 148–164.
- [20] S. Vaswani, L. Lakshmanan, M. Schmidt, et al., Influence maximization with bandits, arXiv preprint arXiv:1503.00024 (2015).
- [21] W. Chen, Y. Wang, Y. Yuan, Q. Wang, Combinatorial multi-armed bandit and its extension to probabilistically triggered arms, The Journal of Machine Learning Research 17 (2016) 1746–1778.
- [22] Z. Wen, B. Kveton, M. Valko, S. Vaswani, Online influence maximization under independent cascade model with semi-bandit feedback, NeurIPS 30 (2017).
- [23] V. Gabillon, B. Kveton, Z. Wen, B. Eriksson, S. Muthukrishnan, Adaptive submodular maximization in bandit setting, NeurIPS 26 (2013).
- [24] D. Das, S. Jain, S. Gujar, Budgeted combinatorial multi-armed bandits, in: AAMAS, 2022, pp. 345–353.
- [25] Y. Azar, D. Panigrahi, N. Touitou, Online graph algorithms with predictions, in: SODA, 2022, pp. 35–66.
- [26] Y. Azar, S. Leonardi, N. Touitou, Flow time scheduling with uncertain processing time, in: STOC, 2021, pp. 1070–1080.
- [27] J. C. Ergun, Z. Feng, S. Silwal, D. Woodruff, S. Zhou, Learning-augmented k-means

clustering, in: ICLR, 2021.

[28] B. Wilder, N. Immorlica, E. Rice, M. Tambe, Maximizing influence in an unknown social network, in: AAAI, 2018, pp. 4743–4750.