

---

# Robust Influence Maximization for Hyperparametric Models

---

Dimitris Kalimeris<sup>1</sup> Gal Kaplun<sup>1</sup> Yaron Singer<sup>1</sup>

## Abstract

In this paper we study the problem of robust influence maximization in the independent cascade model under a hyperparametric assumption. In social networks users influence and are influenced by individuals with similar characteristics and as such they are associated with some features. A recent surging research direction in influence maximization focuses on the case where the edge probabilities on the graph are not arbitrary but are generated as a function of the features of the users and a global hyperparameter. We propose a model where the objective is to maximize the worst-case number of influenced users for any possible value of that hyperparameter. We provide theoretical results showing that proper robust solution in our model is NP-hard and an algorithm that achieves improper robust optimization. We make-use of sampling based techniques and of the renowned multiplicative weight updates algorithm. Additionally we validate our method empirically and prove that it outperforms the state-of-the-art robust influence maximization techniques.

## 1. Introduction

In this paper we study robust influence maximization for hyperparametric diffusion models. First studied by Domingos and Richardson (Domingos & Richardson, 2001) and later elegantly formulated in seminal work by Kempe, Kleinberg, and Tardos (Kempe et al., 2003), influence maximization is the algorithmic task of selecting a small set of individuals who can effectively spread information in a network. The problem we formulate and address in this paper pertains to influence maximization for cases in which the information spread model in the network is subject to some uncertainty.

---

<sup>1</sup>Department of Computer Science, Harvard University, Cambridge, MA, USA. Correspondence to: Dimitris Kalimeris <kalimeris@g.harvard.edu>, Gal Kaplun <galkaplun@g.harvard.edu>.

The most well studied model for information spread is the celebrated Independent Cascade (IC) model. In this model the social network is modeled by a graph and every pair of nodes  $u, v$  that are connected with an edge  $e = (u, v) \in E$  are associated with a probability  $p_e$  that quantifies the probability of  $u$  spreading information to  $v$ . Information spread in this model stochastically progresses from a set of nodes that initiates information to the rest of the nodes in the network as dictated by the graph topology and probabilities encoded on the edges. Influence maximization is then the algorithmic task of selecting a fixed set of individuals that maximize the expected number of nodes that receive information. More formally, given a graph  $G = (V, E, \mathbf{p})$  where  $V$  is the set of nodes,  $E$  is the set of edges, and  $\mathbf{p} \in [0, 1]^{|E|}$  is the vector of edge probabilities, and a parameter  $k \leq |V|$ , influence maximization is the optimization problem:

$$\max_{S: |S| \leq k} f_{\mathbf{p}}(S)$$

where  $f_{\mathbf{p}}(S)$  is the expected number of nodes in the network that receive information when  $S$  is the initial set of nodes that spreads information.

In their seminal work, Kempe, Kleinberg, and Tardos proved that when  $\mathbf{p}$  is known influence maximization can be reduced to monotone submodular maximization under a cardinality constraint. Consequently, a simple greedy algorithm that iteratively selects the node whose marginal influence is approximately maximal obtains a solution that is arbitrarily close to a  $1 - 1/e$  factor of optimal (Nemhauser et al., 1978).

### 1.1. Influence maximization under model uncertainty

In recent years there has been a growing concern regarding the sensitivity of influence maximization to model uncertainty (Goyal et al., 2011; Adiga et al., 2013). In particular, small perturbations or uncertainty regarding the probability vector can have dramatic effects on the quality of a solution (see example illustrated in Figure 1), and even approximating the stability of an instance to small perturbations is intractable (He & Kempe, 2015).

### 1.2. Influence maximization under different models

The sensitivity to small errors that we mentioned in the previous subsection is one issue. The second issue that one needs to consider is that we might actually have several

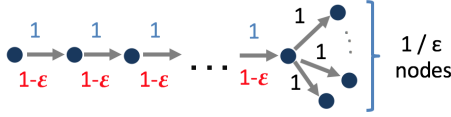


Figure 1. This graph indicates the sensitivity of the influence maximization task to minor differences in the diffusion probabilities and hence the necessity of robust optimization. If all the diffusion probabilities are 1, selecting the node in the beginning of the chain maximizes the influence. However, if the edge probabilities are  $1 - \epsilon$  it will be optimal to select the node in the end of the chain.

models we want to optimize over. For example consider the case where a clothing company wants to advertise shirts and sweaters. The probabilities in the graph will be slightly different but we expect that the set of influential nodes will be more or less the same. Hence, it makes sense to try to identify a set of nodes that are influential, for all the underlying products, i.e. robust for the different models that we care (where each model induces different probabilities).

**Robust influence maximization.** To account for model uncertainty there has recently been a growing body of literature on *robust* influence maximization (He & Kempe, 2016; Chen et al., 2016a; 2017; Ohsaka & Yoshida, 2017; Anari et al., 2019) where the goal is to find a set of nodes whose influence is maximal over an entire set  $\mathcal{P}$  of models:

$$\arg \max_{S: |S| \leq k} \min_{\mathbf{p} \in \mathcal{P}} f_{\mathbf{p}}(S)$$

For a general set of models  $\mathcal{P}$  it is easy to see that the robust influence maximization problem is either trivial or intractable. Specifically, if we have confidence intervals for the diffusion probabilities, i.e.  $p_e \in [c_e^-, c_e^+]$  for all edges  $e$  then  $\max_S \min_{\mathbf{p}} f_{\mathbf{p}}(S)$  simplifies into  $\max_S f_{\mathbf{p}^-}(S)$  where  $\mathbf{p}^- = (c_e^-)_{e \in E}$  due to the monotonicity of  $f_{\mathbf{p}}$  in  $\mathbf{p}$ . In general  $\mathcal{P}$  can be exponentially large hence the problem is intractable. Natural approaches like discretization, sampling, or maximin optimization over the influence function will fail to work for two computational reasons: 1) The space is of exponential dimension in  $|E|$  and 2) The influence function is highly non concave-convex, a form of functions that is amenable to max-min optimization. To circumvent these difficulties, previous work on robust influence maximization have taken two different approaches. The first approach solves the max-min objective but assuming that the number of models is polynomial in the size of the problem (e.g. (Chen et al., 2017; Anari et al., 2019)). The second focuses on the *robust ratio*  $\rho(S) := \min_{\mathbf{p} \in \mathcal{P}} f_{\mathbf{p}}(S) / f_{\mathbf{p}^*}(S)$  where  $S_{\mathbf{p}^*}$  denotes the optimal solution for  $f_{\mathbf{p}}$  (He & Kempe, 2016; Chen et al., 2016a) which is a natural direction that comes with the caveat of not optimizing for the total number of nodes that are influenced.

However, optimizing for the robust ratio is not the right

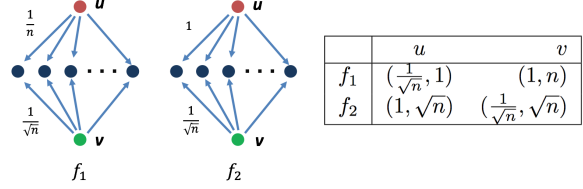


Figure 2. Optimizing for the robust ratio may cost a *multiplicative* factor of  $\sqrt{n}$  in the worst case. Here,  $n$  is the number of blue nodes and we consider  $k = 1$ . In  $f_1$ , the optimal solution is node  $v$  while in  $f_2$  it is  $u$ . The table shows the robust ratios and the expected number of influenced nodes for nodes  $u$  and  $v$  in the form  $(ratio, expected)$ . The ratio objective will choose node  $u$ , influencing only 1 node in expectation. However, the direct approach will select node  $v$  for  $\sqrt{n}$  influenced nodes in expectation. *Maximizing the robust ratio is not the right objective to consider.*

objective to consider and it can be up to a factor of  $\sqrt{n}$  worse than the real robust solution as proved in the following lemma (proof in Appendix A) and illustrated in Figure 2.

**Lemma 1.** *Let  $\mathcal{P}$  be a set of influence functions. Consider the solutions to the objectives:  $\hat{S}_r = \arg \max_{S: |S| \leq k} \rho(S)$  and  $\hat{S}_v = \arg \max_{S: |S| \leq k} \min_{\mathbf{p} \in \mathcal{P}} f_{\mathbf{p}}(S)$ . There exists a set of influence functions  $\mathcal{P}$  for which  $\min_{\mathbf{p}} f_{\mathbf{p}}(\hat{S}_r) = \frac{1}{\sqrt{n}} \min_{\mathbf{p}} f_{\mathbf{p}}(\hat{S}_v)$ , and this approximation ratio is tight.*

**Robust optimization in hyperparameteric models.** A recent line of work in influence maximization and learning in networks explores the interaction of correlations of edge probabilities with the influence spread (Wen et al., 2015; Vaswani et al., 2017; Aral & Dhillon, 2018; Kalimeris et al., 2018). Specifically, it restricts the hypothesis class of IC by imposing correlations on the way that the probabilities are created. It assumes that each node in the network is associated with some features encoding information about it, for example social (age, etc.) or graph-related (degree, pagerank, etc.) characteristics. The influence probability between two nodes is a function of their features and a global *low-dimensional* hyperparameter  $\theta$ . Such approaches have been shown to have sample complexity that only depends on the dimension  $d$  and importantly have been shown to be highly predictive on real information spread data collected from Facebook even with hyperparameteric models parameterized by small number of dimensions (Kalimeris et al., 2018). Intuitively, instead of searching individual probabilities we need to search for a hyperparameter in a much smaller space and once we find the right one it pins down all the influence probabilities.

Low-dimensional hyperparameteric models circumvent the hardness associated with continuous spaces, as they impose structure and the complexity of the influence model is largely determined by the dimension  $d$ . The main question we address in this paper can be informally stated as follows:

*Is there a computationally efficient algorithm to perform robust optimization for hyperparametric models?*

### 1.3. Main result

In this paper we show that in contrast to general influence models, the hyperparametric approach enables tractable solutions to the robust influence maximization problem. At a high level, we show that by a simple sampling procedure one can find an *efficient reduction* from continuous to discrete robust influence maximization that allows us to approximate the value  $\max \min f$  instead of the ratio.

### 1.4. Paper Organization

We begin by formalizing the robust influence maximization problem and the hyperparametric model in Section 2. In Section 3 we describe the main technical result of the paper which allows reducing the continuous robust optimization problem to a discrete problem. In Section 4 we use this reduction and introduce the Hyperparametric Influence Robust Optimizer (HIRO) algorithm for robust influence maximization. In Section 5 we provide a strong hardness result that shows the NP-hardness of robust optimization, even in the sense of bi-criteria approximation. Finally we evaluate the empirical performance of our algorithm in Section 6.

## 2. Preliminaries

A social network is modeled by a graph  $G = (V, E)$  where  $V$  is the set of individuals in the network and  $E$  represents the friendships between them. We use  $n$  and  $m$  to denote  $|V|$  and  $|E|$  respectively. One of the core models for diffusion, the process through which information flows between the nodes of  $G$ , is the Independent Cascade model which was popularized in the seminal work of (Kempe et al., 2003).

**The Independent Cascade (IC) model.** The IC model describes a discrete-step stochastic process through which diffusion spreads from a set of initially active individuals to the rest of the nodes in the network. Each node can be active or inactive and each edge  $e \in E$  in the network is associated with some probability  $p_e$ . All nodes begin as inactive and at time step  $t = 0$  a subset of nodes  $S \subseteq V$ , called the seed set, is chosen and becomes active. At every time step  $t + 1$ , every node  $u$  that became active at time step  $t$  attempts to influence every of its non-active neighbors  $v$ , independently and succeeds with probability  $p_{(u,v)}$ .

**Influence functions.** For a given graph  $G = (V, E)$  and vector of probabilities  $\mathbf{p} \in [0, 1]^m$  the *influence function*  $f_{\mathbf{p}} : 2^V \rightarrow \mathbb{R}$  measures the expected number of nodes that will become influenced in the graph  $G$  for a seed set  $S \subseteq V$ :

$$f_{\mathbf{p}}(S) = \sum_{A \subseteq E} r_A(S) \prod_{e \in A} p_e \prod_{e \notin A} (1 - p_e)$$

where  $r_A(S)$  denotes the number of nodes that are reachable in  $G$  from  $S$  using only edges from  $A$ . An important property of  $f_{\mathbf{p}}$  is that it is *monotone submodular* for any  $\mathbf{p}$ .

**Influence maximization.** For a given influence function  $f_{\mathbf{p}} : 2^{[n]} \rightarrow \mathbb{R}$  and value  $k \leq n$ , influence maximization is the optimization problem:  $\max_{S: |S| \leq k} f_{\mathbf{p}}(S)$ . The problem is NP-hard but since the influence function is monotone and submodular (Kempe et al., 2003) a simple greedy algorithm which iteratively selects nodes whose marginal contribution is largest obtains a solution that is a  $1 - 1/e$  factor of the optimal solution (Nemhauser et al., 1978) and this approximation ratio is optimal unless  $P=NP$  (Feige, 1998).

**Robust influence maximization.** Given a graph  $G = (V, E)$  and a set of probability vectors  $\mathcal{P}$  the goal of *robust* influence maximization is to find a solution of size  $k$  that has high value for every possible influence function  $f_{\mathbf{p}}$  that can be generated by  $\mathbf{p} \in \mathcal{P}$ :

$$\max_{S: |S| \leq k} \min_{\mathbf{p} \in \mathcal{P}} f_{\mathbf{p}}(S)$$

There are two sources of inapproximability known for this problem: first, since it a generalization of influence maximization it is NP-hard to get any approximation better than  $1 - 1/e$  and we are therefore satisfied with solutions that are *approximately* optimal. The second source of inapproximability is due to the fact that our solution space  $\{S : |S| \leq k\}$  is highly non-convex which makes it intractable to obtain *proper* solutions (Figure 3). In particular, it is NP-hard to find a set of size  $k$  that obtains any approximation better than  $\mathcal{O}(\log(n))$  for the robust optimization problem (Chen et al., 2017). For this reason, we seek *bi-criteria approximations*. A solution  $\hat{S}$  is an  $(\alpha, \beta)$  bi-criteria approximation to the max-min solution of size  $k$  if  $\beta |\hat{S}| \leq k$  and:

$$\min_{\mathbf{p} \in \mathcal{P}} f_{\mathbf{p}}(\hat{S}) \geq \alpha \max_{S: |S| \leq k} \min_{\mathbf{p} \in \mathcal{P}} f_{\mathbf{p}}(S)$$

Due to its sources of hardness the gold standard for robust influence maximization are  $(1 - 1/e, \Omega(\log^{-1}(n)))$  bicriteria approximations (Chen et al., 2017; Krause et al., 2008; He & Kempe, 2016; Anari et al., 2019).

**Hyperparametric influence models.** A hyperparametric model  $H : \Theta \times X \rightarrow [0, 1]$  restricts the traditional IC model by imposing correlations between the probabilities of different edges. Each edge  $e = (u, v)$  is associated with a  $d$ -dimensional feature vector  $x_e \in X \subseteq [-1, 1]^d$  encoding information about its endpoints. The probability of  $u$  influencing  $v$  is a function of the features  $x_{(u,v)}$ , and a global hyperparameter  $\theta \in \Theta \subseteq [-B, B]^d$ , for some constant  $B > 0$ . That is:  $p_e = H(\theta, x_e)$ . The most standard hyperparametric models are Generalized Linear Models (GLM) for which:

$$H(\theta, x_e) = h(\theta^\top x_e) + \xi_e$$




Figure 3. The green node denotes the optimal solution in each case. Any deterministic solution has robust influence of 1 (only the selected node). However, the distribution selecting uniformly from  $\{u, v\}$  (improper solution) has expected influence of  $n/2$ .

where  $\xi_e$  is drawn from some bounded distribution. To ease the presentation, throughout this paper we treat the model as if  $\xi_e = 0^1$ . Our results hold for a family of generalized linear models that we define later which includes standard choices for  $h$  are linear, sigmoid, or the logarithm functions.

**Hyperparameter Influence Robust Optimization.** For a given hyperparametric model  $H$ , set of features  $\{x_e\}_{e \in E}$  and hyperparameter space  $\Theta \subseteq [-B, B]^d$  the set of possible diffusion probabilities is:

$$\mathcal{H} := \{ \{H(\theta, x_e)\}_{e \in E} : \theta \in \Theta \}$$

and robust influence maximization then reduces to:

$$\min_{\mathbf{p} \in \mathcal{H}} f_{\mathbf{p}}(\hat{S}) \geq \alpha \max_{S: |S| \leq k} \min_{\mathbf{p} \in \mathcal{H}} f_{\mathbf{p}}(S)$$

Throughout the paper we prove theorems for general feature spaces and hyperparametric models. When the hyperparametric model  $H$  and set of features is clear from context, it will be convenient to use  $f_{\theta}$  instead of  $f_{\mathbf{p}}$  where  $\mathbf{p}$  is the probability vector  $\mathbf{p}$  generated by a hyperparametric model. We will use the abbreviated notation defined above:

$$\min_{\theta \in \Theta} f_{\theta}(\hat{S}) \geq \alpha \cdot \max_{S: |S| \leq k} \min_{\theta \in \Theta} f_{\theta}(S)$$

### 3. Robust Optimization

In this section we describe the main result of the paper. We show that for an extremely broad class of hyperparametric models robust influence maximization is computationally tractable. More specifically, we show that for generalized linear hyperparametric models that are 1-Lipschitz, for any  $\epsilon > 0$ , a natural sampling procedure from the hyperparameter space  $\Theta = [-B, B]^d$  generates  $l \in \tilde{\mathcal{O}}\left(d \left(\frac{n \cdot m \cdot d}{\epsilon}\right)^d \log \frac{1}{\delta}\right)$  influence functions  $f_1, \dots, f_l$  such that with probability  $1 - \delta$ , robust continuous optimization over  $\Theta$  reduces to robust discrete optimization over the functions  $\{f_i\}_{i=1}^l$ . That is, an algorithm that returns  $\hat{S} \subseteq V : |\hat{S}| \leq k$  satisfying:

$$\min_{i \in [l]} f_i(\hat{S}) \geq \alpha \max_{S: |S| \leq k} \min_{i \in [l]} f_i(S)$$

<sup>1</sup>We note that the results carry over when  $\xi_e$  is drawn from a distribution with mean 0 and bounded support.

implies the existence of an algorithm for which:

$$\min_{\theta \in \Theta} f_{\theta}(\hat{S}) \geq \alpha \max_{S: |S| \leq k} \min_{\theta \in \Theta} f_{\theta}(S) - \epsilon.$$

This reduction from the continuous space of infinitely-many functions  $\{f_{\theta}\}_{\theta \in \Theta}$  to polynomially-many functions  $\{f_i\}_{i \in [l]}$  is handled in two steps. We first prove that influence functions that are generated by a class of hyperparametric models that we call *stable* have bounded Lipschitzness (Section 3.1). Using this property we then prove that sampling polynomially-many functions from the hyperparametric model suffices to obtain approximation to the robust objective (Section 3.2). Finally, we show how to produce near optimal solutions to the robust optimization problem defined on  $\{f_{\theta}\}_{\theta \in \Theta}$  by implementing a best-response oracle on a set of sampled functions  $\{f_i\}_{i \in [l]}$  using a Multiplicative Weight Updates (MWU) procedure (Section 4).

#### 3.1. Stability implies Lipschitzness

We now prove that  $f_{\theta}$  is  $L$ -Lipschitz for  $L \in \text{poly}(n)$  if the hyperparametric model that generates it is *stable*.

**Definition 1.** A hyperparametric model  $H : \Theta \times X \rightarrow [0, 1]$  is stable if it is a generalized linear model that is 1-Lipschitz with respect to the  $\ell_1$  norm, i.e. for  $\theta \in \Theta, x_e \in X$  we have that  $H(\theta, x_e) = h(\theta^\top x_e)$  and for every  $\theta, \theta' \in \Theta$ :

$$|h(\theta^\top x_e) - h(\theta'^\top x_e)| \leq \|\theta - \theta'\|_1$$

It is easy to verify that the hyperparametric models used in influence maximization literature are stable: linear  $H(\theta, x_e) = \theta^\top x_e$ , logistic  $H(\theta, x_e) = \frac{1}{1 + \exp(-\theta^\top x_e)}$  and probit  $H(\theta, x_e) = \Phi(\theta^\top x_e)$  where  $\Phi$  is the CDF of the standard Gaussian, appear in (Wen et al., 2015; Vaswani et al., 2017; Kalimeris et al., 2018) and are all stable.

Intuitively, stable hyperparametric models are not sensitive to small changes of the hyperparameter. Hence, despite the fact that a modification on  $\theta$  affects the probabilities in all edges, the difference is not very large and we are able to bound the absolute change in the influence function. The following two lemmas are inspired by (Chen et al., 2016a) and (Chen et al., 2016b) who proved similar results for non-hyperparametric models.

**Lemma 2.** Assume that the hyperparametric model is stable. Then, the influence function  $f_{\theta}$  is Lipschitz with respect to the  $\ell_1$  norm with Lipschitz constant  $L = nm$ .

*Proof.* Since  $H$  is stable we get that for every  $e \in E$  and every two edge probabilities  $p_e, p'_e$  produced by the hyperparametric model by parameters  $\theta, \theta' \in \Theta$ :

$$|p_e - p'_e| = |H(\theta, x_e) - H(\theta', x_e)| \quad (1)$$

$$= |h(\theta^\top x_e) - h(\theta'^\top x_e)| \leq \|\theta - \theta'\|_1. \quad (2)$$

Hence,  $\|\theta - \theta'\|_1 \leq \epsilon$  implies  $|p_e - p'_e| \leq \epsilon$ .

Now notice that the influence function is monotone with respect to the diffusion probabilities. As a result, the maximum change that can occur given the constraint  $|p_e - p'_e| \leq \epsilon$  is when  $p'_e = \min\{p_e + \epsilon, 1\}$  (or  $p'_e = \max\{0, p_e - \epsilon\}$ ) for all  $e \in E$ . We focus on the first case; the second is identical.

Fix a seed set  $S$ . When the probability of an edge  $e = (u, v)$  is increased from  $p_e$  to  $p_e + \epsilon$ , there is an increase in activation probability for all the nodes that are reachable from  $S$  through  $u$  by at most  $\epsilon$ . There are at most  $n$  nodes that are reachable through  $u$  so the total change in the influence of set  $S$  in that case is  $n\epsilon$ . Using the same argument for each edge in the network gives the desired bound of  $nm\epsilon$ .  $\square$

In our method, the Lipschitz parameter is polynomially related to the complexity of implementing a best-response oracle in the MWU procedure. The fact that the Lipschitzness  $L$  is polynomial in the number of nodes in the graph  $n$  is, therefore, a crucial property as it makes the robust optimization problem computationally tractable.

**Tightness of the Lipschitz constant.** The Lipschitzness of the function polynomially determines the computational complexity of our method. As we show, the Lipschitzness bound  $L = n \cdot m$  is tight in the worst case over all possible graphs (proof is deferred to Appendix A).

**Lemma 3.** *There is a graph for which any influence function generated by a stable non-trivial<sup>2</sup> generalized linear model that is continuous has Lipschitzness  $L = nm$ .*

### 3.2. Covering $\Theta$ via sampling

The Lipschitzness property from Lemma 2 is useful since it bounds the change in the estimation of influence function  $f_\theta$  when we approximate it with another function  $f_{\theta'}$  for  $\theta, \theta'$  that are  $\epsilon$ -close, i.e.  $\|\theta - \theta'\|_1 \leq \epsilon$ . Thus, if we can construct a set of representative parameters  $\Theta_\epsilon$  s.t.  $\forall \theta \in \Theta$  there is some  $\theta'$  which is  $\epsilon$  close with respect to the  $\ell_1$  norm we can reduce the robust optimization problem over the infinite hyperparameter space to the finite case. We know now this can be done using a reasonable number of samples from  $\Theta$ .

**Lemma 4.** *Let  $\Theta = [-B, B]^d$  and  $\epsilon, \delta > 0$ . If we sample a set  $\Theta_\epsilon$  of size  $s \in \tilde{\mathcal{O}}\left(d\left(\frac{Bd}{\epsilon}\right)^d \log \frac{1}{\delta}\right)$  uniformly at random from  $\Theta$ , then with probability at least  $1 - \delta$ , for any  $\theta \in \Theta$  there exists  $\theta' \in \Theta_\epsilon$  such that  $\|\theta - \theta'\|_1 \leq \epsilon$ .*

*Proof.* It suffices to partition the hypercube  $[-B, B]^d$  into  $\ell_1$  balls of radius  $\epsilon$  and bound the number of points necessary to sample in order to have at least one point in each  $\ell_1$  ball. Since every point in  $\Theta$  is covered by a ball and every two

<sup>2</sup> $h : \mathbb{R} \rightarrow [\lambda, 1 - \lambda]$  is non-trivial if it is continuous and surjective for some  $x_e \in X$ .

points in the ball are close, sampling points that correspond to every ball is a representative set of parameters that are  $\epsilon$ -close w.r.t. the  $\ell_1$  distance to any parameter in  $\Theta$ .

It is well known that  $r = \left(\frac{2B \cdot d}{\epsilon}\right)^d$  balls of  $\ell_1$  radius  $\epsilon$  suffice to cover the  $d$ -dimensional hypercube  $[-B, B]^d$ . Let  $b_1, \dots, b_r$  denote the covering balls and  $\theta_1, \dots, \theta_s$  be points drawn uniformly at random from  $\Theta = [-B, B]^d$ . Then:

$$\begin{aligned} \mathbb{P}[\exists b_j \text{ such that } \nexists \theta_i \in b_j] &\leq \sum_j \mathbb{P}[\nexists \theta_i \in b_j] \\ &\leq r \mathbb{P}[\nexists \theta_i \in b_1] \\ &\leq r \left(1 - \frac{1}{r}\right)^s \end{aligned}$$

Using  $s \in \mathcal{O}\left(d\left(\frac{Bd}{\epsilon}\right)^d \log \frac{Bd}{\epsilon\delta}\right)$  samples we get:

$$\mathbb{P}[\exists b_j \text{ such that } \nexists \theta_i \in b_j] \leq \delta$$

Thus, since  $\Theta \subseteq \cup_j b_j$  and the radius of the balls is  $\epsilon$  we get that for every  $\theta \in \Theta$  there exists a  $\theta' \in \{\theta_i\}_{i \in [s]}$  that is  $\epsilon$  close w.r.t.  $\ell_1$  norm with probability at least  $1 - \delta$ .  $\square$

Furthermore, we can extend the notion of *covering* a convex space  $\Theta$  to the coverage of a family of functions.

**Definition 2.** *Let  $\mathcal{F} = \{f_\theta : 2^V \rightarrow \mathbb{R} \mid \theta \in \Theta\}$  be a family of influence functions and  $\mathcal{F}_\epsilon \subset \mathcal{F}$  such that  $|\mathcal{F}_\epsilon| < \infty$ . We say that  $\mathcal{F}_\epsilon$   $\epsilon$ -covers  $\mathcal{F}$  if for any  $f \in \mathcal{F}$  there exists an  $f_\epsilon \in \mathcal{F}_\epsilon$  and any  $S \subseteq V$  s.t.  $|f(S) - f_\epsilon(S)| \leq \epsilon$ .*

The following corollary is obtained by fusing lemmas 2 and 4 with Definition 2.

**Corollary 1.** *Let  $\mathcal{F} = \{f_\theta : 2^V \rightarrow \mathbb{R} \mid \theta \in \Theta\}$  be the family of influence functions and let  $\mathcal{F}_\epsilon$  be sampled uniformly at random from  $\mathcal{F}$ , such that  $|\mathcal{F}_\epsilon| \in \tilde{\mathcal{O}}\left(d\left(\frac{LBd}{\epsilon}\right)^d \log \frac{1}{\delta}\right)$ . Then  $\mathcal{F}_\epsilon$   $\epsilon$ -covers  $\mathcal{F}$  with probability at least  $1 - \delta$ .*

## 4. Reducing Continuous to Discrete RO

In the previous section we proved that since the influence function is Lipschitz, the infinite family of functions  $\mathcal{F} = \{f_\theta \mid \theta \in \Theta\}$  can be well-approximated by the finite (and crucially, polynomially-sized) family  $\mathcal{F}_\epsilon = \{f_{\theta_\epsilon} \mid \theta_\epsilon \in \Theta_\epsilon\}$ , where  $\Theta_\epsilon$  is obtained by sampling u.a.r. from  $\Theta$ . Hence, the task of robust influence maximization in the hyperparametric model intuitively reduces to finding a procedure to perform robust optimization on a finite set of functions.

The following lemma formalizes this intuition by stating that, under mild conditions,  $\alpha$ -approximate continuous robust optimization reduces to  $\alpha$ -approximate discrete robust optimization. The proof is deferred to Appendix A.

**Lemma 5.** *Let  $\mathcal{F} = \{f_\theta : 2^V \rightarrow \mathbb{R} \mid \theta \in \Theta \subseteq [-B, B]^d\}$  be a family of influence functions. Consider a family*

**Algorithm 1** HIRO: Hyperparam Inf Robust Optimizer

**Input:**  $G = (V, E), \{x_e\}_{e \in E}, H : \Theta \times X \rightarrow [0, 1], \epsilon, \delta$

$$l \leftarrow \tilde{\mathcal{O}} \left( d \left( \frac{LBd}{\epsilon} \right)^d \log \frac{1}{\delta} \right)$$

$$T \leftarrow \tilde{\mathcal{O}} \left( \frac{d(\log n + \log \log \frac{1}{\delta})}{\epsilon^2} \right), \eta \leftarrow \frac{\log l}{2T}$$

$f_1, \dots, f_l \leftarrow \text{SAMPLE}(G, H, \Theta, \{x_e\}_{e \in E})$   
 $w_1[1], \dots, w_1[l] \leftarrow 1/l, \dots, 1/l$

**for** each time step  $t \in [T]$  **do**

$$S_t \leftarrow \text{GREEDY} \left( \sum_{i=1}^l w_t[i] f_i \right)$$

**for** each  $i \in [l]$  **do**

$$w_t[i] \propto \exp \left\{ -\eta \sum_{\tau=1}^{t-1} f_i(S_\tau) \right\}$$

**end for**

**end for**

**Output:** select  $S$  u.a.r. from  $\{S_1, S_2, \dots, S_T\}$

$\mathcal{F}_\epsilon \subset \mathcal{F}$  s.t.  $\mathcal{F}_\epsilon$   $\epsilon$ -covers  $\mathcal{F}$ . Then,  $\alpha$ -approximate robust optimization on  $\mathcal{F}$  reduces to  $\alpha$ -approximate robust optimization on  $\mathcal{F}_\epsilon$ . That is, an algorithm that returns  $\hat{S} \subseteq V$ :

$$\min_{\hat{S} \in \mathcal{F}_\epsilon} f_\theta(\hat{S}) \geq \alpha \cdot \max_{S \subseteq V} \min_{f_\theta \in \mathcal{F}_\epsilon} f_\theta(S)$$

implies an algorithm that for any  $\epsilon > 0$  returns  $\hat{S} \subseteq V$  s.t.:

$$\min_{\hat{S} \in \mathcal{F}} f_\theta(\hat{S}) \geq \alpha \cdot \max_{S \subseteq V} \min_{f_\theta \in \mathcal{F}} f_\theta(S) - \epsilon.$$

**The HIRO algorithm.** Given the reduction from the continuous hyperparametric problem to the discrete we can now describe the Hyperparametric Influence Robust Optimizer (HIRO) Algorithm 1 which gives an optimal bi-criteria approximation to the robust influence maximization in the hyperparametric setting. HIRO takes as input a graph  $G$ , the edge features  $\{x_e\}_{e \in E}$ , the hyperparametric model  $(H, \Theta)$  that dictates the edge probabilities and an error parameter  $\epsilon$  that controls the quality of the returned solution. The output is a set of  $k$  nodes. It starts by sampling  $l$  points from  $\Theta$  and hence, constructing  $l$  different influence functions that serve as a proxy for the continuous problem. It assigns uniform weights to these functions and runs MWU for a number of steps that depends on  $\epsilon$ , where in each step higher emphasis is placed on the ones with poor historical performance. In every iteration it optimizes a convex combination of the functions, which is possible since all  $f_i$ s are monotone submodular, and hence amenable to optimization using the GREEDY algorithm. It keeps the outcome as a candidate solution. In the end, one of the candidate solutions is returned u.a.r. The intuition is that since each solution is good for some iteration of the algorithm (meaning for some specific weighting of the  $f_i$ s), on expectation the solution

that is returned is good for all the functions that performed poorly for some iteration of the MWU and hence, robust.

**Theorem 1.** HIRO with error parameter  $\epsilon$ , runs in time  $\text{poly}(n, \epsilon, \log(1/\delta))$  and returns the uniform distribution  $\mathcal{U}$  over solutions  $\{S_1, \dots, S_T\}$ , s.t. with probability at least  $1 - \delta$ :

$$\min_{\theta \in \Theta} \mathbb{E}_{S \sim \mathcal{U}} [f_\theta(S)] \geq \left(1 - \frac{1}{e}\right) \max_{S: |S| \leq k} \min_{\theta \in \Theta} f_\theta(S) - 2\epsilon.$$

where  $T \in \tilde{\mathcal{O}} \left( \frac{d(\log n + \log \log \frac{1}{\delta})}{\epsilon^2} \right)$ .

*Proof.* From Corollary 1 we know that constructing  $l$  different influence functions  $f_1, \dots, f_l$  by sampling  $l = \tilde{\mathcal{O}} \left( d \left( \frac{LBd}{\epsilon} \right)^d \log \frac{1}{\delta} \right)$  points u.a.r. from  $\Theta$  yields an  $\epsilon$ -cover of  $\Theta$  with probability at least  $1 - \delta$ .

Since  $L \leq n \cdot m$  (Lemma 2) and  $d$  is constant according to the main assumption of the hyperparametric model, we know that constructing the cover takes polynomial time.

(Chen et al., 2017) proved that for influence functions  $f_1, \dots, f_l$ , the MWU procedure, run for  $T$  iterations, with a learning rate of  $\eta = \log(l)/2T$  that uses GREEDY as an approximate best-response oracle, returns the uniform distribution over  $T$  solutions s.t.:

$$\min_{i \in [l]} \mathbb{E}_{\hat{S}} [f_i(\hat{S})] \geq \left(1 - \frac{1}{e}\right) \max_S \min_{i \in [l]} f_i(S) - \mathcal{O} \left( \sqrt{\frac{\log(l)}{T}} \right)$$

By applying Lemma 5 we get:

$$\min_{\theta \in \Theta} \mathbb{E}_{S \sim \mathcal{U}} [f_\theta(\hat{S})] \geq \left(1 - \frac{1}{e}\right) \max_S \min_{\theta \in \Theta} f_\theta(S) - \epsilon - \mathcal{O} \left( \sqrt{\frac{\log(l)}{T}} \right)$$

For  $l \in \tilde{\mathcal{O}} \left( d \left( \frac{LBd}{\epsilon} \right)^d \log \frac{1}{\delta} \right)$  we get the desired bound by setting  $T \in \tilde{\mathcal{O}} \left( \frac{d(\log n + \log \log \frac{1}{\delta})}{\epsilon^2} \right)$  as required, since  $L \cdot B \in \mathcal{O}(n)$ .  $\square$

The result of the previous theorem implies the following bi-criteria approximation guarantee by returning  $\hat{S} = \cup_{t=1}^T S_t$  instead of the uniform distribution over the  $\{S_1, \dots, S_T\}$ .

**Corollary 2.** HIRO with error parameter  $\epsilon$ , returns a set  $\hat{S}$  of size at most  $\tilde{\mathcal{O}} \left( \frac{d(\log n + \log \log(1/\delta))}{\epsilon^2} k \right)$  which, with probability at least  $1 - \delta$ , satisfies:

$$\min_{\theta \in \Theta} f_\theta(\hat{S}) \geq \left(1 - \frac{1}{e}\right) \max_{S: |S| \leq k} \min_{\theta \in \Theta} f_\theta(S) - 2\epsilon.$$

## 5. Lower Bound

In this section, we prove that a structural assumption such as the hyperparametric restriction of the IC model is vital for

robust influence maximization, since otherwise we might need to sample a number of functions that is exponential in  $n$ . We do so by providing a strong hardness result: define the family of functions  $\mathcal{F}_{\mathbf{p}} = \{f_{\mathbf{p}} \mid \mathbf{p} \in \mathcal{P}\}$  for some finite  $\mathcal{P}$ . We prove that the problem  $\max_{S: |S| \leq k} \min_{\mathbf{p} \in \mathcal{P}} f_{\mathbf{p}}(S)$  is NP-hard to approximate within any constant factor with a reasonable bicriteria approximation.

This is in sharp contrast with the main result about hyperparametric robust influence maximization that we proved in Theorem 1 and Corollary 2 since, assuming the hyperparametric model,  $\mathcal{P}$  is always of polynomial size and hence we only need to increase our budget by a factor of  $\tilde{O}(\log n)$ . In Theorem 2 we formalize the impossibility result. We provide a proof sketch, the full proof is in Appendix B.

**Theorem 2.** *Let  $G$  be a graph with  $n$  nodes and  $m$  edges, and  $\delta, \epsilon > 0$ . There is no algorithm to find a set  $\hat{S}$  of size  $|\hat{S}| \leq (1 - \delta) \ln |\mathcal{P}| \cdot k$  that achieves an approximation factor better than  $\mathcal{O}(\frac{1}{n^{1-\epsilon}})$  to the problem  $\max_{S: |S| \leq k} \min_{\mathbf{p} \in \mathcal{P}} f_{\mathbf{p}}(S)$ , where  $\mathcal{P} \subseteq \{\lambda, 1 - \lambda\}^m$ ,  $\lambda = o(1)$ , and  $|\mathcal{P}| \in \Omega(\text{poly}(n))$ , unless  $P = NP$ .*

This essentially means that if in our problem, we need to construct a cover that contains more than a polynomial number of functions then, the best bicriteria approximation we can hope for will contain significantly more than  $k$  nodes.

*Proof.* Our reduction is based on (He & Kempe, 2016). The authors there prove the hardness of a different version of robust influence maximization where the objective is to approximate well the individual optima for a set of different influence functions instead of influencing as many nodes as possible in the worst case. We reduce from GAP SET COVER. The use of the gap version of the problem is to show that even if we augment the budget of nodes by a factor of roughly  $\log |\mathcal{P}|$  the problem remains NP-hard.

Given an instance of gap set cover, we construct a bipartite graph on  $n$  nodes,  $m$  edges, and  $\text{poly}(n)$  different probability sets that correspond to the different influence functions, such that when there is a set cover of size at most  $k$  then there exists a seed set  $S$  for which all the influence functions have high value, while when there is no set cover of size  $\log n \cdot k$ , then there is at least one influence function that has low value for any seed set, even if we allow sets of size  $\log n \cdot k$ . The asymptotic difference between the values of the objective functions, with and without the cover set, enables the decision of the gap set cover.  $\square$

## 6. Experiments

To measure the empirical performance of the robust hyperparametric approach we conducted four sets of experiments: (1) First, we examine how many functions are empirically required for the sufficient covering of  $\mathcal{F} = \{f_{\theta} \mid \theta \in \Theta\}$

and observe that in practice it is far smaller than the theoretical worst-case upper bound of Corollary 1. (2) We examine the rate of convergence of HIRO to the robust solution and show substantially faster convergence than in Theorem 1. (3) We benchmark the performance of HIRO against other methods for robust influence maximization and observed that it consistently outperforms all previous methods. (4) We measure how well a seed set of size  $k$  performs compared to a seed set of slightly augmented budget according to the bi-criteria approximation guarantee (Corollary 2).

**Graphs.** We generated four different synthetic networks using standard random graph models to analyze the impact of topological variations among different social networks. All networks were generated with  $n = 500$  vertices. We used the Barabási-Albert preferential attachment model where each new node is connected to 4 (preferably high-degree) existing ones, the Watts-Strogatz small world model where each node is connected to 5 nodes in the ring topology and the probability of rewiring an edge is  $3/n$ , Erdős-Rényi random graph model with edge-construction probability  $p = 3/n$  and the configuration model with a power law degree distribution and  $\alpha = 2$ . For a more detailed description of these models please refer to Appendix C.

**Hyperparametric model.** We used the sigmoid function as the hyperparameteric model to determine the diffusion probabilities, i.e.  $h(\theta^{\top} x_e) = \frac{1}{1 + \exp(-\theta^{\top} x_e)}$  as in (Kalimeris et al., 2018). We generated  $d$  random features in  $[-1, 1]$  for every edge. We used  $d = 5$ , however our results are consistent across a large range of dimensions  $d$  and feature-generating techniques, such as normal or uniform distributions over the unit hyper-cube  $[-1, 1]^d$  and it’s discrete analog  $\{-1, 1\}^d$ . We sampled  $\Theta_{\epsilon} = \{\theta_1, \dots, \theta_l\}$  from  $\Theta = [-1, 1]^d$  and generated the family of influence functions  $\mathcal{F}_{\epsilon} = \{f_i \mid \theta_i \in \Theta_{\epsilon}\}$  for  $l = 20$ . In addition we set  $T = 10$  HIRO iterations with the exception of Experiments 1 and 2, where  $l$  and  $T$  are the free variable, respectively.

**Benchmarks.** We benchmark the performance of HIRO with respect to the following algorithms:

- *Random Seed:* We select  $k$  nodes u.a.r.; to account for variance, we average over the solutions from 100 trials.
- *Top  $k$  Degree nodes:* In this benchmark we chose the  $k$  highest-degree nodes in the graph.
- *Random Greedy:* In robust optimization problems, a typical approach is to randomize over the best strategies against any possible influence function. Consequently, this method runs GREEDY algorithm on every function  $f_i$  for  $\theta_i \in \Theta_{\epsilon}$  and chooses one of the outputs uniformly at random.
- *LUGREEDY:* We compare ourselves against Algorithm 2 in (Chen et al., 2016a) and refer to it as LUGREEDY. LUGREEDY is an algorithm specifically developed for robust influence maximization.

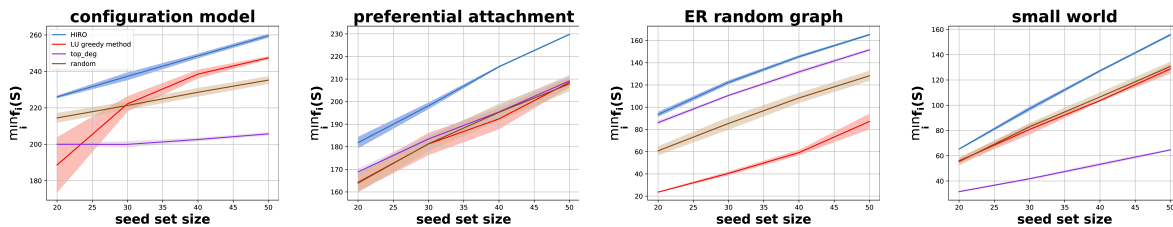


Figure 4. Comparison to benchmark solutions of the robust problem.




Figure 5. Illustration of the gap between a regular seed set size and the bi-criteria augmented budget.

It is oblivious to the structure of the model and only accounts for the confidence interval of each edge. LUGREEDY produces two solutions, by running the GREEDY algorithm twice - over both the lower and upper boundaries of the confidence intervals. Then, the algorithm chooses the best solution out of the two, assuming the lower boundaries as true probabilities. LUGREEDY achieves the best approximation for the robust ratio objective, mentioned in the introduction.

## 6.1. Experimental results

We perform 50 trials for each experiment and plot mean and standard deviation in the figures.

**Experiment 1.** (Figure 6, Appendix C) The goal of this experiment is to understand the sample complexity of the hyperparameter, i.e. how many functions do we need to sample to approximate the robust solution accurately. We sample  $l = 50$  values from  $\Theta$  as a benchmark (can be seen as a test/validation set). In that way we create a cover  $\mathcal{F}_\varepsilon$  that serves as a proxy for  $\mathcal{F} = \{f_\theta | \theta \in \Theta\}$ . We plot  $\min_{1 \leq i \leq l} f_i(S_r)$  where  $S_r$  is the result of running HIRO over  $r \in \{1, 10, 20, 30, 40, 50\}$  functions. We plot three such trials with different seed sizes  $k \in \{10, 25, 50\}$ . Figure 6, deferred to the appendix, demonstrates that in practice sample complexity plateaus, even though theoretically increase in the number of functions should bring about an increase in the value of the robust solution.

**Experiment 2.** (Figure 7, Appendix C) We show that HIRO converges rapidly in practice, that is, even after a constant number of iterations, HIRO achieves great results. We run the experiment three times, for different values

of  $k \in \{10, 25, 50\}$ . For each trial, we run HIRO for a  $T \in \{1, 5, 10, 15\}$  number of iterations. Figure 7, found in the appendix, illustrates this inquiry. Notwithstanding the small world model, which has slow growth, other plots indeed seem to converge quickly for all values of  $k$ .

**Experiment 3.** (Figure 4) This experiment compares HIRO with the benchmarks and illustrates consistency of our algorithm. We plot the value of a benchmark as a function of the size of the seed set. We can see that across all generative models and proposed benchmarks, even after few iterations, our algorithm performs at the very top, with the lowest variance. It is readily seen that other benchmarks are competitive when the seed set is small, but as the seed set grows so does the gap in performance of the best heuristic.

**Experiment 4.** (Figure 5) We evaluate the gap in the value of the robust solution between a seed set of size  $k$  and a seed set of size  $\beta \cdot k \log n$ . Recall that the HIRO algorithm chooses at random among  $T$  solutions a seed with size  $k$ . Instead, we take a union of these solutions and return a size- $k'$  subset of the union, where  $\beta \in (0, 1]$ . Here we choose  $k = 10$  and we report the results in Figure 5.

## 7. Conclusion

In this paper, we proposed a new formulation of robust influence maximization by utilizing a very broad class of hyperparametric models. We provided an efficient reduction from continuous to discrete robust influence maximization and an optimal and computationally tractable algorithm for the problem in terms of bi-criteria approximation. We empirically assessed its performance and found that it consistently surpasses state-of-the-art methods.



## References

- Adiga, A., Kuhlman, C., Mortveit, H., and Kumar, A. Sensitivity of diffusion dynamics to network uncertainty. In *AAAI Conf. on Artificial Intelligence*, 2013.
- Anari, N., Haghtalab, N., Naor, J. S., Pokutta, S., Singh, M., and Torricok, A. Structured robust submodular maximization: offline and online algorithms. In *The 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- Aral, L. and Dhillon, P. Social influence maximization under empirical influence models. In *Nature Human Behavior*, pp. 375–382, 2018.
- Chen, R., Lucier, B., Singer, Y., and Syrgkanis, V. Robust optimization for non-convex objectives. In *Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
- Chen, W., Lin, T., Tan, Z., Zhao, M., and Zhou, X. Robust influence maximization. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016a.
- Chen, W., Wang, Y., and Yuan, Y. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. In *Journal of Machine Learning Research*, volume 17, pp. 1–33, 2016b.
- Dinur, I. and Steurer, D. Analytical approach to parallel repetition. In *In Proceedings 45th ACM Symposium on Theory of Computing (STOC)*, pp. 624–633, 2014.
- Domingos, P. and Richardson, M. Mining the network value of customers. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 57–66. ACM, 2001.
- Feige, U. A threshold of  $\ln n$  for approximating set cover. In *Journal of the ACM (JACM)*, volume 45(4), pp. 634–652, 1998.
- Goyal, A., Bonchi, F., and Lakshmanan, V. A data-based approach to social influence maximization. In *Proc. VLDB Endowment*, 5(1), pp. 73–84, 2011.
- He, X. and Kempe, D. Stability of influence maximization. *CoRR*, abs/1501.04579, 2015. URL <http://arxiv.org/abs/1501.04579>.
- He, X. and Kempe, D. Robust influence maximization. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016.
- Kalimeris, D., Singer, Y., Subbian, K., and Weinsberg, U. Learning diffusion using hyperparameters. In *International Conference on Machine Learning (ICML)*, 2018.
- Karp, R. M. Reducibility among combinatorial problems. In *R. E. Miller and J. W. Thatcher (editors). Complexity of Computer Computations. New York: Plenum.*, pp. 85–103, 1972.
- Kempe, D., Kleinberg, J., and Tardos, É. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146. ACM, 2003.
- Krause, A., McMahan, H. B., Guestrin, C., and Gupta, A. Robust submodular observation selection. In *Journal of Machine Learning Research*, volume 9, pp. 2761–2801, 2008.
- Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. An analysis of approximations for maximizing submodular set functions. In *Mathematical Programming 14(1)*, pp. 265–294, 1978.
- Ohsaka, N. and Yoshida, Y. Portfolio optimization for influence spread. In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3-7, 2017*, pp. 977–985, 2017. doi: 10.1145/3038912.3052628. URL <https://doi.org/10.1145/3038912.3052628>.
- Vaswani, S., Kveton, B., Wen, Z., Ghavamzadeh, M., Lakshmanan, L. V. S., and Schmidt, M. Model-independent online learning for influence maximization. In *International Conference on Machine Learning (ICML)*, 2017.
- Wen, Z., Kveton, B., and Ashkan, A. Efficient learning in large-scale combinatorial semi-bandits. In *International Conference on Machine Learning (ICML 2015)*, pp. 1113–1122, 2015.