

Sample Complexity Bounds for Influence Maximization

Gal Sadeh

Tel Aviv University, Israel
galsdh@gmail.com

Edith Cohen 

Google Research, Mountain View, CA, USA
Tel Aviv University, Israel
edith@cohenwang.com

Haim Kaplan

Google Research, Tel Aviv, Israel
Tel Aviv University, Israel
haimk@tau.ac.il

Abstract

Influence maximization (IM) is the problem of finding for a given $s \geq 1$ a set S of $|S| = s$ nodes in a network with maximum influence. With stochastic diffusion models, the influence of a set S of seed nodes is defined as the expectation of its *reachability* over *simulations*, where each simulation specifies a deterministic reachability function. Two well-studied special cases are the *Independent Cascade (IC)* and the *Linear Threshold (LT)* models of Kempe, Kleinberg, and Tardos [31]. The influence function in stochastic diffusion is unbiasedly estimated by averaging reachability values over i.i.d. simulations. We study the IM *sample complexity*: the number of simulations needed to determine a $(1 - \epsilon)$ -approximate maximizer with confidence $1 - \delta$. Our main result is a surprising upper bound of $O(s\tau\epsilon^{-2} \ln \frac{n}{\delta})$ for a broad class of models that includes IC and LT models and their mixtures, where n is the number of nodes and τ is the number of diffusion steps. Generally $\tau \ll n$, so this significantly improves over the generic upper bound of $O(sn\epsilon^{-2} \ln \frac{n}{\delta})$. Our sample complexity bounds are derived from novel upper bounds on the variance of the reachability that allow for small relative error for influential sets and additive error when influence is small. Moreover, we provide a data-adaptive method that can detect and utilize fewer simulations on models where it suffices. Finally, we provide an efficient greedy design that computes an $(1 - 1/e - \epsilon)$ -approximate maximizer from simulations and applies to any submodular stochastic diffusion model that satisfies the variance bounds.

2012 ACM Subject Classification Theory of computation \rightarrow Machine learning theory

Keywords and phrases Sample complexity, Influence maximization, Submodular maximization

Digital Object Identifier 10.4230/LIPIcs.ITCS.2020.29

Funding This research is partially supported by the Israel Science Foundation (Grant No. 1841/14 and Grant No. 1595/19).

1 Introduction

Models for the spread of information among networked entities were studied for decades in sociology and economics [26, 20, 29]. A diffusion process is initiated from a seed set of nodes (entities) and progresses in steps: Initially, only the seed nodes are activated. In each step additional nodes may become active based on the current set of active nodes. The progression can be deterministic or stochastic. The t -stepped influence of a seed set S of nodes is then defined as its expected reachability (total number of active nodes) in t steps.

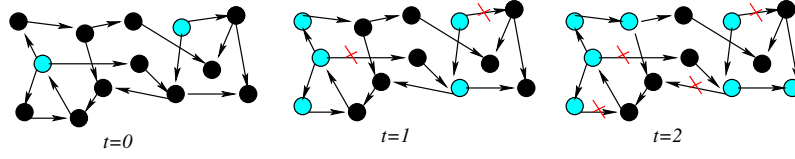


© Gal Sadeh, Edith Cohen, and Haim Kaplan;
licensed under Creative Commons License CC-BY
11th Innovations in Theoretical Computer Science Conference (ITCS 2020).

Editor: Thomas Vidick; Article No. 29; pp. 29:1–29:36



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** A 2-step diffusion from two seed nodes in a live-edge model.

Influence maximization (IM) is the problem of finding a set S of nodes of specified cardinality $|S| = s$ and maximum *influence*. The IM problem was formulated nearly two decades ago by Richardson and Domingos [17, 40] and inspired by the application of viral marketing. In a seminal paper, Kempe, Kleinberg, and Tardos [31] studied stochastic diffusion models and introduced two elegant special cases, the *Independent Cascade* (IC) and *Generalized Threshold* (GT) diffusion models. Their work sparked extensive followup research and large scale implementations [35, 7, 30, 39]. Currently IM is applied in multiple domains with linked entities for tasks as varied as diversity-maximization (the most representative subset of the population) and sensor placement that maximize coverage [32, 4, 34].

We consider *stochastic diffusion models* (SDM) $\mathcal{G}(V)$ over $|V| = n$ nodes that are specified by a distribution $\phi \sim \mathcal{G}$ over sets $\phi := \{\phi_v\}_{v \in V}$ of monotone non-decreasing boolean *activation functions*

$$\phi_v : 2^{V \setminus \{v\}} \rightarrow \{0, 1\}.$$

A diffusion process starts with a seed set $S \subset V$ of nodes and $\phi \sim \mathcal{G}$. At step 0 we activate the seed nodes $\text{Reach}^0(\phi, S) := S$. The diffusion then proceeds deterministically: At step $t > 0$ all active nodes remain active and we activate any inactive node v where $\phi_v(\text{Reach}^{t-1}(\phi, S)) = 1$:

$$\text{Reach}^{t+1}(\phi, S) := \{v \in V \mid \phi_v(\text{Reach}^t(\phi, S)) = 1\}.$$

The τ -steps *reachability set* of a seed set S is the random variable $\text{Reach}^\tau(\phi, S)$ for $\phi \sim \mathcal{G}$ and respectively the τ -steps *reachability*, $R^\tau(S)$, is the random variable that is the number of active nodes $|\text{Reach}^\tau(\phi, S)|$ for $\phi \sim \mathcal{G}$. Finally, the influence value of S is defined to be the expectation

$$I^\tau(S) := \mathbb{E}[R^\tau(S)] = \mathbb{E}_{\phi \sim \mathcal{G}}[|\text{Reach}^\tau(\phi, S)|].$$

We refer to the case where the diffusion is allowed to progress until there is no growth as *unrestricted* diffusion and this corresponds to $\tau = n - 1$. The influence $I^\tau(S)$ is a monotone set function. We say that an SDM is *submodular* when the influence function is submodular and that it is *independent* if the activation functions ϕ_v of different nodes are independent random variables. The IM problem for seed set size s and τ steps is to find

$$\arg \max_{S: |S| \leq k} I^\tau(S).$$

The reader might be more familiar with well-studied special cases of this general formulation. *Live-edge* diffusion models $\mathcal{G}(V, \mathcal{E})$ are specified by a graph (V, \mathcal{E}) with $|V| = n$ nodes and $|\mathcal{E}| = m$ directed edges and a distribution $E \sim \mathcal{G}$ over subsets $E \subset \mathcal{E}$ of “live” edges. When expressed as an SDM, the activation functions that correspond to E have $\phi_v(T) = 1$ if and only if there is an edge from a node in T to v in the graph (V, E) . Live-edge models are always submodular: This because $|\text{Reach}^\tau(E, S)|$, which is the number of nodes reachable

from S in (V, E) by paths of length at most τ , is a coverage function and hence monotone and submodular. Therefore, so is the influence function $I^\tau(S)$, which is an expectation of a distribution over coverage functions. A live-edge model is independent if we only have dependencies between incoming edges to the same node. The Independent Cascade (IC) model is the special case of an independent live-edge model where all edges $e \in \mathcal{E}$ are independent Bernoulli random variables selected with probabilities p_e ($e \in \mathcal{E}$).

Another well-studied class are *generalized threshold* (GT) models [31, 35]. A GT model $\mathcal{G}(V, \mathbf{f})$ is specified by a set $\mathbf{f} := \{f_v\}_{v \in V}$ of monotone functions $f_v : 2^V \rightarrow [0, 1]$. The randomization is specified by a set of threshold values $\theta \sim \mathcal{G}$ where $\theta := \{\theta_v\}_{v \in V}$. The corresponding activation functions to θ are

$$\phi_v(T) := \text{Indicator}(\theta_v \leq f_v(T)).$$

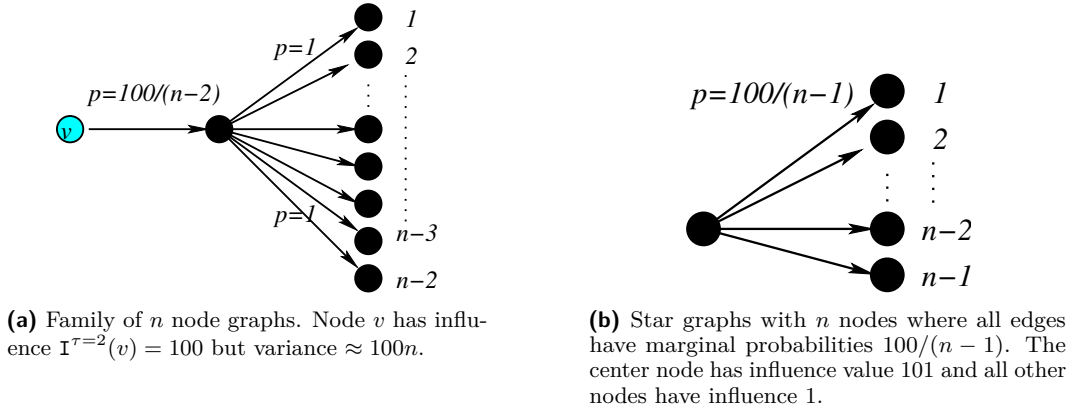
An important subclass are *Independent GT* (IGT) where we require that \mathbf{f} are submodular and nodes $v \in V$ have independent threshold values $\theta_v \sim U[0, 1]$. Mossel and Roch [35, 36] proved that IGT models are submodular, which is surprising because the functions $|\text{Reach}^\tau(\phi, S)|$ are generally not submodular. Their proof was provided for unrestricted diffusion but extends to the case where we stop the process after τ steps. Finally, Linear threshold (LT) models [26, 31] are a special case of IGT where we have an underlying directed graph and each edge (u, v) is associated with a fixed weight value $b_{uv} \geq 0$ so that for all $v \in V$ $\sum_u b_{uv} \leq 1$ and the functions are defined as the sums $f_v(A) := \sum_{u \in A \cap N(v)} b_{uv}$. Kempe et al showed [31] that each LT model is equivalent to an independent live-edge model.

One of the challenges brought on by the IM formulation is computational efficiency. Kempe et al [31] noted that the IM problem generalizes the classic Max Cover problem even with $\tau = 1$ and a live-edge model with a fixed set of live edges ($p_e = 1$ for all $e \in \mathcal{E}$). Therefore, IM inherits Max Cover's hardness of approximation for ratio better than $1 - (1 - 1/s)^s \geq 1 - 1/e$ [21] for a cover with s sets. On the positive, with submodular models, an approximation ratio of $1 - (1 - 1/s)^s$ can be achieved by the first s nodes of a greedy sequence generated by sequentially adding a node with maximum marginal value [37]. A challenge of applying Greedy with stochastic models, however, is that even point-wise evaluation of the influence function can be computationally intensive. Exact evaluation even for IC models [6] and LT models [8] is #P hard. As for approximation, Kempe et al proposed to work with *averaging oracles*

$$\hat{A}^\tau(T) := \frac{1}{\ell} \sum_{i=1}^{\ell} |\text{Reach}^\tau(\phi_i, T)|$$

that average the reachability values obtained from a set $\{\phi_i\}_{i=1}^{\ell}$ of i.i.d. simulations. Recall that in the general SDM formulation, a simulation is specified by a set ϕ of node activation functions. For live-edge models, a simulation is simply a set of concurrently live edges E . In GT models, a simulation is specified by a set of thresholds θ .

Averaging oracles, which we focus on here, have several appealing properties: First, it is robust compared to estimators tailored to models that satisfy specific assumptions (see related work section) in that for any diffusion model \mathcal{G} , also with complex and unknown dependencies (between activation functions of different nodes or between edges in live-edge models), for any set S , $\hat{A}(S)$ is an unbiased estimate of the exact influence value $I^\tau(S)$ and estimates are accurate as long as the variance of $R^\tau(S)$ is "sufficiently" small. Second, the oracle is constructed directly from simulations and does not require learning or inferring the underlying diffusion model that generated the data [41, 24, 23, 18]. Therefore, the results are not sensitive to modeling assumptions and learning accuracy [9, 27]. Moreover, estimation of



■ **Figure 2** Example live-edge models.

model parameters even for IC models can require a large number of simulations: Consider the family of IC models depicted in the example in Figure 2b where tiny edge probabilities that require many simulations to estimate are critical for IM accuracy. Finally, when the reachability functions $\text{Reach}^\tau(\phi, T)$ are monotone and submodular (as is the case with live-edge models), so is their average \hat{A} , and hence the oracle optimum can be approximated by the greedy algorithm. Prior work addressed the efficiency of working with averaging oracles by improving the efficiency of greedy maximization [32, 25] and applied sketches [10] to efficiently estimate $\hat{A}(S)$ values [7, 14].

The fundamental question we study here is the *sample complexity* of IM, that is, the number of i.i.d. simulations needed to recover an approximate maximizer of the influence function I^τ . Formally, for parameters (ϵ, δ) , identify a seed set T of size s so that $\Pr[I^\tau(T) \geq (1 - \epsilon)\text{OPT}_s^\tau] \geq 1 - \delta$, where $\text{OPT}_s^\tau := \max_{S \mid |S| \leq s} I^\tau(S)$ is the exact maximum. Note that the recovery itself is generally computationally hard and the sample complexity only considers the information we can glean from a set of simulations.

Kempe et al provided an upper bound of

$$O\left(\epsilon^{-2} s n \log \frac{n}{\delta}\right) \quad (1)$$

on the sample complexity of the harder *Uniform Relative-Error Estimation (UREE)* problem where for a given (ϵ, δ) we bound the number of simulations so that with probability $1 - \delta$, for all subsets S such that $|S| \leq s$, $\hat{A}(S)$ approximates $I^\tau(S)$ within relative error of ϵ . The sample complexity of UREE upper bounds that of IM because the oracle maximizer $\arg \max_{S \mid |S| \leq s} \hat{A}(S)$ must be an approximate maximizer. We provide the argument for the bound (1) here because it is basic and broadly applies to all SDMs: The reachability values $\text{Reach}^\tau(\phi, S)$, and hence their expectation, $I^\tau(S)$ have values in $[1, n]$. Using the multiplicative Chernoff bound (with values divided by n) we obtain that $O(\epsilon^{-2} n \ln \delta^{-1})$ simulations guarantee a relative error of ϵ with probability at least $(1 - \delta)$ for the estimate of any particular set S . Interestingly, this sample complexity bound is tight for point-wise influence estimation even for IC models: The example family of models in Figure 2a is such that $\tau = 2$ and $\Omega(\epsilon^{-2} n)$ simulations are required for estimating the influence value of a single node. The UREE sample complexity bound (1) follows from applying a union bound over all $\binom{n}{s} = O(n^s)$ subsets.

The generic upper bound has prohibitive linear dependence on the number of nodes n (that the example in Figure 2a shows is unavoidable for UREE even for IC models). A simple example shows that we can not hope for an umbrella improvement for IM: Consider the star

graph family of Figure 2b when edges are dependent so that either all edges are live or none is. Clearly $n/100$ simulations are necessary to detect a 1-step approximate maximizer (which must be the actual maximizer). The remaining hope is that we can obtain stronger bounds on the IM sample complexity for models with weaker or no dependencies such as IC and IGT models. This question eluded researchers for nearly two decades.

1.1 Contributions and Overview

We study the sample complexity of influence maximization from averaging oracles computed from i.i.d. simulations. One of our main contributions is an upper bound of

$$O\left(\epsilon^{-2} s \tau \log \frac{n}{\delta}\right) \quad (2)$$

on the IM sample complexity of independent strongly submodular SDMs. Informally, strong submodularity means that the influence function of any “reduced” model (model derived from original one by setting a subset $T \subset V$ of nodes as active) is submodular. The IC and IGT models are special cases of strongly submodular independent SDMs.

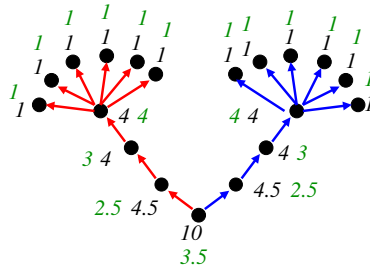
Interestingly, we provide similar sample complexity bounds for natural families of models that are not independent: Mixtures of small number of strongly submodular SDMs and what we call *b-dependence* live-edge models that allow for positive dependence of small groups of edges with a shared tail node.

Our bound improves over prior work by replacing the prohibitive linear dependence in the number of nodes n in (1) with the typically much smaller value τ . While on worst-case instances unrestricted diffusions may require $\Omega(n)$ steps, understanding the sample complexity in terms of τ is important: First, IM with explicit step limits [5, 33, 22, 19, 15], is studied for applications where activation time matters. Moreover, due to the “small world” phenomenon [44], in “natural” settings we can expect most activations (even with unrestricted diffusions) to occur within a small number of steps. In the latter case, unrestricted influence values are approximated well by corresponding step-limited influence with $\tau \ll n$.

Our improvement is surprising as generally a linear-in- n number of simulations is necessary for estimating influence values of some nodes or to estimate essential model parameters (for example, the edge probabilities in IC models), and this is the case even when τ is very small. This shows that the maximization problem is in an information sense inherently easier and can circumvent these barriers.

We overview our results and implications – complete proofs can be found in the appendix. We review related work in Section 2 and place it in the context of our results. In Section 3 we formulate quality measures for influence oracles and relate unrestricted and step-limited influence. In particular, we observe that for IM it suffices that the oracle provides good estimates (within a small relative error) of larger influence values. This allows us to circumvent the lower bound for point-wise relative-error estimates shown in Example 2a.

In Section 4 we state our main technical result that for independent strongly submodular SDMs, upper bounds the variance of the reachability of set T $\text{Var}[\mathbf{R}^\tau(T)]$ by $\tau \mathbf{I}^\tau(T) \max_{v \in V \setminus T} \mathbf{I}^{\tau-1}(v)$. This variance upper bound facilitates estimates with small relative error for sets with larger influence values and additive error for sets with small influence values. We also provide a family of IC models that shows that the linear dependence on τ in the variance bound is necessary. We derive similar variance bounds to mixtures of strongly submodular independent SDMs and *b-dependence* models. All our subsequent sample complexity bounds apply generically to any SDM (submodular/independent or not) that satisfies variance bounds of this form. In Section 5 we review *averaging oracles* and bound the sample



■ **Figure 3** Dependent SDM Example. A live-edge model where with probability $\frac{1}{2}$ all red edges are active and otherwise all blue edges are active. The influence values $\Gamma^4(v)$ are shown in black. Simulation averages and RR samples with full simulations provide unbiased estimates of influence values $E[\hat{\Gamma}^4(v)] = \Gamma^4(v)$. However, “efficient” RRS, which works with the marginal edge probabilities ($p_e = \frac{1}{2}$) or with decomposed simulations is biased, with $E[\hat{\Gamma}^4(v)]$ shown in green. We can see that the bias induces large errors and also yields an erroneous maximizer.

complexity using variance upper bounds. In section 6 we present our *median-of-averages* oracle that amplifies the confidence guarantees of the averaging oracle and facilitates a tighter sample complexity bound. In Section 7 we provide a data-adaptive framework that provides guarantees while avoiding the worst-case sample complexity upper bounds on models when a smaller number of simulations suffices. In Section 8 we consider computational efficiency and present a greedy maximization algorithm based on our median-of-averages oracles that returns a $(1 - (1 - 1/s)^s - \epsilon)$ approximate maximizer with probability $1 - \delta$. The design generically applies to any SDM with a submodular influence function that satisfies the variance bounds.

2 Related work

We focus here on influence estimates obtained from averages of i.i.d. simulations of a SDM. We note that alternative approaches can be more effective for specific families of models. Most notably for IC models, state of the art large-scale greedy approximate maximization algorithms [43, 42, 38, 28] are not based on simulation averages. The estimates are also obtained by building a sample that reflects the influence of each but instead they use a finer building block of i.i.d. *Reverse Reachability (RR) searches*. The random RR search method was proposed in [10] to estimate size of reachability sets in graphs and Borg et al [3] adapted it to IC models.

The method is applicable in principle with any live-edge model: A basic RR search is conducted by selecting a node $v \in V$ uniformly at random and performing a BFS search on reversed edges that is pruned at length τ . The search “flips” edges as their head node is reached according to conditional distribution on \mathcal{G} . The index number of the RR search is then added to the sample set of each node that is “reached” by the search. Influence of a subset S can then be unbiasedly estimated from the cardinality of the union of the samples of nodes $v \in S$ and the greedy algorithm can be applied to the sets of samples for approximate maximization. To obtain an approximate influence maximizer we need to perform RR searches until some node has a sample of size $O(\epsilon^{-2} s \log(n/\delta))$. In the worst case, this requires $O(\epsilon^{-2} s n \log(n/\delta))$ RR searches. For general live-edge models, an independent RR search can always be obtained from a simulation $E \sim \mathcal{G}$ by randomly drawing a node and performing a reverse search from it using edges E . The same simulation, however, can not

generally be reused to generate multiple independent RR searches. This way of obtaining RR searches works for general live-edge models (with arbitrary dependencies) but requires $O(\epsilon^{-2}ns \log(n/\delta))$ simulations, which does not improve over the generic upper bound (1).

The appeal of the RR searches method is that it can be implemented very efficiently for independent live-edge (including IC or LT) models. The total work performed requires only $O(\epsilon^{-2}ms \log(n/\delta))$ “edge flips” that can be easily performed using specified edge probabilities p_e for IC models. Moreover, the basic building block of RR searches are local simulations of sets of incoming edges of specified nodes and the full computation requires at most $O(\epsilon^{-2}s \log(n/\delta))$ local simulations for each node. When we have full simulations generated by an independent live-edge model these “local” simulations are independent and the required number of “local simulations” can be obtained by decomposing $O(\epsilon^{-2}s \log(n/\delta))$ full simulations. But the caveat is that this approach breaks the coherence of simulations, as we construct each RR search from components taken from multiple simulations. These “efficient” implementations (i.e. based on decomposed simulations or edge flips according to marginal probabilities) may “catastrophically fail” when dependencies exist: The influence estimates obtained are biased and cause large errors even when the variance is low. Figure 3 shows an example of a simple mixture model (of two degenerate IC models) where “efficient” RRS has large error due to bias but averages of few simulations provide accurate estimates. To summarize, with RRS, the implementation that works with full simulations is robust to dependencies but is inefficient and the efficient implementation breaks ungracefully even with light dependencies. To conclude, both basic approaches to approximate IM, simulation averages and RRS offer distinct advantages: Simulation averages are robust in that they remain unbiased and are accurate on any SDM, including dependent ones, for which the variance is sufficiently small whereas RRS offers more efficiency for purely independent live-edge models.

3 Preliminaries

We consider stochastic diffusion models $\mathcal{G}(V)$ as outlined in the introduction. We denote by $\text{Reach}^\tau(\phi, T)$ the τ -steps reachability set of T when we use a specific set ϕ of activation functions. We will use the notation $\text{Reach}^\tau(T)$ (with the parameter ϕ omitted) for the random variable $\text{Reach}^\tau(\phi, T)$ obtained when we draw $\phi \sim \mathcal{G}$ according to the model.

3.1 Utility Functions

For simplicity, the discussion in the introduction took the utility of a reachable set to be the number of reachable nodes $\mathbb{V}\text{Reach}^\tau(\phi, T) := |\text{Reach}^\tau(\phi, T)|$. Generally, we can consider *utility functions* $H : 2^V \rightarrow \mathbb{R}_+$ that are nonnegative monotone non-decreasing with $H(\emptyset) = 0$:

$$\mathbb{V}\text{Reach}^\tau(\phi, T) := H(\text{Reach}^\tau(\phi, T)). \quad (3)$$

Submodular utility is particularly natural and studied by Mossel and Roch [35]. Additive utility is the special case where nodes have nonnegative weights $w : V \rightarrow \mathbb{R}^+$ and

$$\mathbb{V}\text{Reach}^\tau(\phi, T) := \sum_{v \in \text{Reach}^\tau(\phi, T)} w(v). \quad (4)$$

We consider a diffusion model $\mathcal{G}(V, H)$ together with a utility function H . The random variable $\mathbb{R}_{\mathcal{G}}^{\tau}(T)$ is the utility of the reachable set, that is, $\mathbb{V}\text{Reach}^{\tau}(\phi, T)$ when $\phi \sim \mathcal{G}$. The influence function is then the expected utility of the reachable set

$$\mathbb{I}^{\tau}(T) := \mathbb{E}[\mathbb{R}^{\tau}(T)] = \mathbb{E}_{\phi \sim \mathcal{G}} \mathbb{V}\text{Reach}^{\tau}(\phi, T).$$

We denote the maximum influence value of a subset of cardinality s by $\text{OPT}_s^{\tau} := \max_{S: |S| \leq s} \mathbb{I}^{\tau}(S)$.

It follows from the definition that for any SDM $\mathcal{G}(V, H)$ with utility H , the influence $\mathbb{I}^{\tau}(T)$ is monotone non-decreasing in τ and in the set T and the optimum values OPT_s^{τ} are non-decreasing in τ and s . Generally, influence functions $\mathbb{I}^{\tau}(T)$ of SDMs may not be submodular even when utility is additive. The influence function is known to be submodular for live-edge and for IGT models [35] with submodular utility.

3.2 Reduced Models

We work with the following notion of model reduction. Let $\mathcal{G}(V, H)$ be an independent SDM with submodular utility. For a set of nodes $T \subset V$, we define the *reduced model* $\mathcal{G}'(V', H')$ of \mathcal{G} with respect to T : The reduced model contains the nodes $V' = V \setminus T$. The activation function $\phi'_v \sim \mathcal{G}'$ for $v \in V \setminus T$ are obtained by drawing $\phi_v \sim \mathcal{G}$ conditioned on $\phi_v(T) = 0$ and take

$$\text{for all } S \subset V \setminus (T \cup \{v\}), \phi'_v(S) := \phi_v(S \cup T)$$

(Note that since we have independent SDM we can separately consider the distribution of activation functions of each node). The utility in \mathcal{G}' is the marginal utility in \mathcal{G} with respect to T :

$$\text{for all } S \subset V \setminus T, H'(S) := H(S \cup T) - H(T).$$

The reduced model $\mathcal{G}'(V', H')$ is also an independent SDM with submodular utility: Activation functions $\phi' \sim \mathcal{G}'$ are independent and monotone and the utility is monotone with $H'(\emptyset) = 0$ and submodular.

3.3 Strongly Submodular SDM

We say that an independent SDM $\mathcal{G}(V, H)$ is *strongly submodular* if the utility function H is submodular and the influence function $\mathbb{I}_{\mathcal{G}}^{\tau}$ is submodular with any reduced model \mathcal{G}' and step limit $\tau \geq 0$. IC and IGT models are strongly submodular SDMs (see Theorem 19).

The variance and thus sample complexity upper bounds that we present in the sequel apply to any strongly submodular SDM. We will also provide bounds for some dependent families of models. One family is a slight generalization of IC models that we refer to as *b-dependence*. Here edges are partitioned into disjoint groups, where each group contains at most b edges emanating from the same node. The edges in a group must be either all live or none live (are positively dependent).

3.4 Relating Step-Limited and Unrestricted Influence

When unrestricted diffusion from a seed set S is such that most activations occur within τ steps, the unrestricted influence $\mathbb{I}(S)$ is approximated well by τ -step influence $\mathbb{I}^{\tau}(S)$. We can also relate unrestricted influence with small expected steps-to-activation to step-limited

influence: For a seed set S , node v , and length d , we denote by $p(S, v, d)$ the probability that node v is activated in a diffusion from S in step d . For additive utility functions (4) by definition, $\mathbf{I}^\tau(S) = \sum_{v \in V} w(v) \sum_{d \leq \tau} p(S, v, d)$. The *expected length of an activation path* from S (in unrestricted diffusion) is:

$$\overline{D}(S) := \frac{\sum_{v \in V} w(v) \sum_{d \leq \tau} d \cdot p(S, v, d)}{\mathbf{I}(S)}. \quad (5)$$

The following lemma is an immediate consequence of Markov's inequality and shows that τ -stepped influence with $\tau = O(\overline{D}(S))$ approximates well the unrestricted influence:

► **Lemma 1.** *For all S and $\epsilon > 0$, $\mathbf{I}^{\overline{D}(S)\epsilon^{-1}}(S) \geq (1 - \epsilon)\mathbf{I}(S)$.*

3.5 Influence Oracles

We say that a set function \hat{F} is an ϵ -approximation of another set function F at a point T if $|\hat{F}(T) - F(T)| \leq \epsilon \max\{F(T), \text{OPT}_1(F)\}$, where $\text{OPT}_s(F) := \max_{S' \mid |S'| \leq s} F(S')$. That is, the estimate \hat{F} has a small relative error for sets T with $F(T) \geq \text{OPT}_1(F)$ and a small absolute error of $\epsilon \text{OPT}_1(F)$ for sets T with $F(T) \leq \text{OPT}_1(F)$. We say that \hat{F} provides a *uniform ϵ -approximation* for all subsets T in a collection \mathcal{C} if \hat{F} is an ϵ -approximation for all $T \in \mathcal{C}$.

An *influence oracle*, $\hat{\mathbf{I}}^\tau$, is a randomized data structure that is constructed from a set of i.i.d. simulations of a model. The influence oracle, $\hat{\mathbf{I}}^\tau$, defines a set function (we use the same name $\hat{\mathbf{I}}^\tau$ for the set function) that for any input query set $T \subset V$, returns a value $\hat{\mathbf{I}}^\tau(T)$. For $\epsilon < 1$ and $\delta < 1$ we say that an oracle provides (ϵ, δ) *approximation guarantees with respect to \mathbf{I}^τ* if for any set T it is an ϵ -approximation with probability at least $1 - \delta$. That is

$$\forall T \text{ such that } \mathbf{I}^\tau(T) \geq \text{OPT}_1^\tau, \Pr \left[\frac{|\hat{\mathbf{I}}^\tau(T) - \mathbf{I}^\tau(T)|}{\mathbf{I}^\tau(T)} \geq \epsilon \right] \leq \delta. \quad (6)$$

$$\forall T \text{ such that } \mathbf{I}^\tau(T) \leq \text{OPT}_1^\tau, \Pr \left[|\hat{\mathbf{I}}^\tau(T) - \mathbf{I}^\tau(T)| \geq \epsilon \text{OPT}_1^\tau \right] \leq \delta. \quad (7)$$

where $\text{OPT}_1^\tau := \text{OPT}_1(\mathbf{I}^\tau)$. Example 2a shows that this type of requirement is what we can hope for with an oracle that is constructed from a small number of simulations.

The (ϵ, δ) requirements are *for each* particular set T . If we are interested in stronger guarantees that with probability $(1 - \delta)$ the approximation uniformly holds *for all* sets in a collection \mathcal{C} , we can use an oracle that provides $(\epsilon, \delta_A = \delta/|\mathcal{C}|)$ guarantees. The ϵ -approximation guarantee for all sets in \mathcal{C} then follow using a union bound argument: The probability that all $|\mathcal{C}|$ sets are approximated correctly is at most $|\mathcal{C}|\delta_A \leq \delta$.

4 Variance Bounds

We consider upper bounds on the variance $\text{Var}[\mathbf{R}^\tau(T)]$ of the reachability of a set of nodes T that have the following particular form

$$\text{Var}[\mathbf{R}^\tau(T)] \leq c\mathbf{I}^\tau(T) \max\{\mathbf{I}^\tau(T), \max_{v \in V} \mathbf{I}^\tau(v)\} \quad (8)$$

for some $c \geq 1$. The sample complexity bounds we present in the sequel apply to any SDM that satisfies these bounds. In the remaining part of this section we state our variance upper bound for strongly submodular SDMs and extensions and a tight worst-case lower bound for IC models.

4.1 Variance Upper Bound

The following key theorem facilitates our main results. We show that any strongly submodular SDM satisfy the bound (8) with $c = \tau$. The proof is technical and provided in Appendix A.

► **Theorem 2** (Variance Upper Bound Lemma). *Let $\mathcal{G}(V, H)$ be a strongly submodular SDM. Then for any step limit $\tau \geq 0$ and a set $T \subset V$ of nodes we have*

$$\text{Var}[\mathbf{R}^\tau(T)] \leq \tau \mathbf{I}^\tau(T) \max_{v \in V \setminus T} \mathbf{I}^{\tau-1}(v).$$

Some natural dependent SDMs have a variance bound of the form (8): (See Appendix F for proofs.)

► **Corollary 3.** *IC models with τ -steps and b -dependence satisfy the bound (8) with $c = 2b\tau$. Any mixture of τ -steps strongly submodular SDMs where each model has probability at least p satisfy the bound (8) with $c = (\tau + 1)/p$.*

4.2 Variance Lower Bound

We provide a family of IC models for which this variance upper bound is asymptotically tight. This shows that the dependence of the variance bound on τ is necessary.

► **Theorem 4** (Variance Lower Bound). *For any $\tau > 0$ there is an IC model $\mathcal{G}^\tau = (V, \mathcal{E})$ with a node $v \in V$ of maximum influence such that $\text{Var}[\mathbf{R}^\tau(v)] \geq \frac{1}{12}\tau \mathbf{I}^\tau(v)^2$.*

Our family of models $\mathcal{G}^\tau = (V, \mathcal{E})$ are such that (V, \mathcal{E}) is a complete directed binary tree of depth $\tau \geq 1$ rooted at $v \in V$ with all edges directed away from the root and $p_e = 1/2$ for all $e \in \mathcal{E}$. We show (details in Appendix B) that:

$$\begin{aligned} \mathbf{I}^\tau(v) &= \tau \\ \text{Var}[\mathbf{R}^\tau(v)] &= \frac{1}{12}\tau(\tau-1)(2\tau-1). \end{aligned}$$

5 The Averaging Oracle

The *averaging oracle* uses i.i.d. simulations $\{\phi_i\}_{i=1}^\ell$. For a query T it returns the average utility of the reachability set of T : $\hat{\mathbf{A}}^\tau(T) = \text{Ave}_{i \in [\ell]} \mathbf{VReach}^\tau(\phi_i, T) := \frac{1}{\ell} \sum_{i=1}^\ell \mathbf{VReach}^\tau(\phi_i, T)$. We quantify the approximation guarantees of an averaging oracle in terms of a variance bound of the form (8).

► **Lemma 5.** *Consider an SDM that for some $c \geq 1$ satisfies a variance bound of the form (8). Then for any $\epsilon, \delta < 1$, an averaging oracle constructed from $\ell \geq \epsilon^{-2}\delta^{-1}c$ i.i.d. simulations provides (ϵ, δ) guarantees.*

In particular for strongly submodular SDMs, we use the variance bound in Theorem 2 and obtain these approximation guarantees using $\ell \geq \epsilon^{-2}\delta^{-1}\tau$ i.i.d. simulations.

Proof. Using variance properties of the average of i.i.d. random variables, we get that for any query T

$$\text{Var}[\hat{\mathbf{A}}^\tau(T)] = \frac{1}{\ell} \text{Var}[\mathbf{R}^\tau(T)] \leq \frac{1}{\ell} c \mathbf{I}^\tau(T) \max\{\mathbf{I}^\tau(T), \text{OPT}_1^\tau\}.$$

The claims follow using Chebyshev's inequality that states that for any random variable X and M , $\Pr[|X - \mathbf{E}[X]| \geq \epsilon M] \leq \epsilon^{-2} \text{Var}[X]/M^2$. We apply it to the random variable $\hat{\mathbf{A}}^\tau(T)$ that has expectation $\mathbf{I}^\tau(T)$ and plug in the variance bound. To establish (6) we use $M = \mathbf{I}^\tau(T)$ and to establish (7) we use $M = \text{OPT}_1^\tau$. ◀

5.1 Sketched Averaging Oracle

For live-edge models with additive utility (4), the query efficiency of the averaging oracle can be improved with off-the-shelf use of τ -step combined reachability sketches [10, 14, 15, 11]. The sketching is according to a sketch-size parameter k that also determines the sketches computation time and accuracy of the estimates that sketches provide. A sketch of size $O(k)$ is computed for each node v so that for any set of nodes S , $\sum_{i=1}^r \text{VReach}^\tau(E_i, S)$ can be efficiently estimated from the sketches of the nodes $v \in S$. The computation of the sketches from an arbitrary set of simulations $\{E_i\}$ uses at most $\sum_i |E_i| + k \sum_v \max_i d_v(E_i)$ edge traversals, where $d_v(E_i)$ is the maximum in-degree of node v over simulations $\{E_i\}$. In the case of an IC model, the expected number of traversals is $(k + \ell) \sum_e p_e$. Sketching with general node weights can be handled as in [11]. The estimates obtained from the sketches are unbiased with coefficient of variation $1/\sqrt{k-2}$ and are concentrated: Sketches of size $k = O(\epsilon^{-2} \log(\delta^{-1}))$ provide estimates with relative error ϵ with probability $1 - \delta$.

6 Confidence Amplification: The Median-of-Averages Oracle

The statistical guarantees we provide for our averaging oracle are derived from variance bounds. The limitation is that the number of simulations we need to provide (ϵ, δ) guarantees is linear in δ^{-1} and therefore the number of simulations we need to provide uniform guarantees (via a union bound argument) grows linearly with the number of subsets. In order to find an approximate optimizer, we would like to have a uniform ϵ -approximation for all the $\binom{n}{s}$ subsets of size at most s but doing so with an averaging oracle would require too many simulations. We adapt to our setting a classic confidence amplification technique [1] to construct an oracle where the number of simulations grows logarithmically in the confidence parameter δ^{-1} .

A *median-of-averages* oracle is specified by a number r of *pools* with ℓ simulations in each pool. The oracle is therefore constructed from $r\ell$ i.i.d. simulations ϕ_{ij} for $i \in [r]$ and $j \in [\ell]$.

The simulations of each pool are used in an averaging oracle that for the i th pool ($i \in [r]$) returns the estimates $\hat{A}_i^\tau(T)$. The median-of-averages oracle returns the median value of these r estimates

$$\widehat{\text{mA}}^\tau(T) := \text{Median}_{i \in [r]} \hat{A}_i^\tau(T) = \text{Median Ave}_{i \in [r] \ j \in [\ell]} \text{VReach}^\tau(\phi_{ij}, T). \quad (9)$$

We establish that when the i.i.d simulations are from a model that has variance bound (8) for some $c \geq 1$, the median-of-averages oracle provides (ϵ, δ) approximation guarantees using $112\epsilon^{-2}c \ln \delta^{-1}$ i.i.d. simulations.

► **Lemma 6.** *Consider an SDM that for some $c \geq 1$ satisfies the variance bound (8). Then for every ϵ and δ , a median-average oracle $\widehat{\text{mA}}$ organized with $r = 28 \ln \delta^{-1}$ pools of $\ell = 4\epsilon^{-2}c$ simulations in each provides (ϵ, δ) approximation guarantees.*

Proof. An averaging oracle with ℓ simulations provides (ϵ, δ_A) approximation guarantees for $\delta_A = 1/4$. Therefore, the probability of correct estimate for any subset is at least $3/4$. We now consider the estimates \hat{A}_j obtained from the r pools when sorted in increasing order. The estimates that are not correct (too low or too high) will be at the prefix and suffix of the sorted order. The expected number of correct estimates is at least $\mu \geq \frac{3}{4}r$. The probability that the median estimate is not correct is bounded by the probability that number of correct estimates is $\leq r/2$, which is $\leq \frac{2}{3}\mu$. From multiplicative Chernoff bounds, the probability of a sum of Bernoulli random variables being below $(1 - \epsilon')\mu$ is at most $e^{-\epsilon'^2\mu/(2+\epsilon')}$. Using $\epsilon' = 1/3$ we have $\epsilon'^2\mu/(2 + \epsilon') = \frac{1}{9} \frac{3}{4} \frac{3}{7} 28 \ln \delta^{-1} = \ln \delta^{-1}$. ◀

29:12 Sample Complexity Bounds for Influence Maximization

As a corollary, we obtain a sample complexity bound for influence maximization from variance bounds:

► **Theorem 7.** *Consider an SDM that satisfies the variance bound (8) for some $c \geq 1$. Then for any $\epsilon < 1$ and $\delta < 1$, using $112\epsilon^{-2}cs \ln \frac{n}{\delta}$ i.i.d. simulations, we can construct a median-of-averages oracle $\widehat{m\mathbf{A}}$ such that the oracle optimum*

$$T := \arg \max_{S \mid |S| \leq s} \widehat{m\mathbf{A}}(S)$$

satisfies

$$\Pr [I^\tau(T) \geq (1 - 2\epsilon)\text{OPT}_s^\tau] \geq 1 - \delta .$$

Proof. We construct our median-of-averages oracle with $\ell = 4\epsilon^{-2}c$ and $r = 28 \ln \delta_{MA}^{-1}$ where $\delta_{MA} = \delta / \binom{n}{s}$. From Lemma 6 using a union bound over the $\binom{n}{s}$ sets we obtain that with probability $1 - \delta$ the oracle provides a uniform ϵ -approximation for all subsets of size at most s . Let S be a set with maximum influence $I(S) = \text{OPT}_s^\tau$ and let T be the oracle optimum. We have

$$I(T) \geq (1 - \epsilon)\widehat{m\mathbf{A}}(T) \geq (1 - \epsilon)\widehat{m\mathbf{A}}(S) \geq (1 - \epsilon)^2 I(S) \geq (1 - 2\epsilon)\text{OPT}_s^\tau .$$

We comment that the $(1 - 2\epsilon)$ ratio is not tight and we can obtain a bound closer to $(1 - \epsilon)$. This because the particular set S to be approximated more tightly by the oracle (that uses enough simulations to support a union bound). ◀

7 Optimization with Adaptive Sample Size

The bound on the number of simulations we derived in Theorem 7 (through a median-of-averages oracle) and also the naive bound (1) (for the averaging oracle) are worst-case. This is obtained by using enough simulations to have the oracle provide a uniform ϵ -approximation with probability at least $1 - \delta$ on any problem instance. To obtain the uniform approximation we applied a union bound over $\binom{n}{s}$ subsets that resulted in an increase in the number of required simulations by an $s \log n$ factor over the base (ϵ, δ) approximation guarantees.

On real data sets a much smaller number of simulations than this worst-case often suffices. We are interested in algorithms that adapt to such data and return a seed set of approximate maximum influence using a respectively smaller number of simulations and while providing statistical guarantees on the quality of the end result. To do so, we apply an adaptive optimization framework [12] (some example applications are [14, 38, 16, 13]). This framework consists of a “wrapper” that take as inputs oracle constructions from simulations and a base algorithm that performs an optimization over an oracle. The wrapper invokes the algorithm on oracles constructed using an increasing number of simulations until a validation condition on the quality of the result is met. The details are provided in Appendix E. We denote by $r(\epsilon, \delta)$ the number of simulations that provides (ϵ, δ) guarantees and we obtain the following results:

► **Theorem 8.** *Suppose that on our data the averaging (respectively, median-of-averages) oracle \hat{I} has the property that with r simulations, with probability at least $1 - \delta$, the oracle optimum $T := \arg \max_{S \mid |S| \leq s} \hat{I}(S)$ satisfies*

$$I^\tau(T) \geq (1 - \epsilon)\text{OPT}_s^\tau .$$

Then with probability at least $1 - 5\delta$, when using

$$2 \max\{r, r(\epsilon, \delta)\} + O\left(\epsilon^{-2}c \left(\ln \frac{1}{\delta} + \ln \left(\ln \ln \frac{n}{\delta} + \ln s\right)\right)\right)$$

simulations with the median-of-averages oracle and

$$2 \max\{r, r(\epsilon, \delta)\} + O\left(\epsilon^{-2}c \left(\ln \frac{1}{\delta} + \ln \left(\ln \ln \frac{n}{\delta} + \ln n\right)\right)\right)$$

simulations with the averaging oracle, the wrapper outputs a set T such that $I^r(T) \geq (1 - 5\epsilon)\text{OPT}_s^r$.

The wrapper can also be used with a base algorithm that is an approximation algorithm. For live-edge models, our averaging oracle is monotone and submodular and hence we can apply greedy to efficiently compute a set with approximation ratio at least $1 - 1/e$ (with respect to the oracle). If we use greedy as our base algorithm we obtain the following:

► **Theorem 9.** *If the averaging oracle \hat{A} is submodular and has the property that with $\geq r$ simulations, with probability at least $1 - \delta$, it provides a uniform ϵ -approximation for all subsets of size at most s , then with*

$$2 \max\{r, r(\epsilon, \delta)\} + O\left(\epsilon^{-2}c \left(\ln \frac{1}{\delta} + \ln \left(\ln \ln \frac{n}{\delta} + \ln n\right)\right)\right)$$

simulations we can find in polynomial time a $(1 - (1 - 1/s)^s)(1 - 5\epsilon)$ approximate solution with confidence $1 - 5\delta$.

8 Approximate Greedy Maximization

In this section we consider the computational efficiency of maximization over our oracle \hat{I} that approximates a monotone submodular influence function I^r . The maximization problem is computationally hard: The brute force method evaluates $\hat{I}(S)$ on all $\binom{n}{s}$ subsets S of size s in order to find the oracle maximizer. An efficient algorithm for approximate maximization of a monotone submodular function \hat{F} is greedy that sequentially builds a seed set S by adding a node u with maximum marginal contribution $\arg \max_{u \in V} (\hat{F}(S \cup \{u\}) - \hat{F}(S))$ at each step. To implement greedy we only need to evaluate at each step the function on a linear number of subsets $\hat{F}(S \cup \{u\})$ for $u \in V$ and thus overall we do sn evaluations of \hat{F} on subsets. With a monotone and submodular \hat{F} , for any $s \geq 1$ the subset T that consists of the first s nodes in a greedy sequence satisfies [37]:

$$\hat{F}(T) \geq (1 - (1 - 1/s)^s) \max_{S \mid |S| \leq s} \hat{F}(S) \geq (1 - 1/e)\text{OPT}_s(\hat{F}).$$

If our functions \hat{F} provides a uniform ϵ -approximation of another function F for all subsets of size at most s , then $F(T) \geq (1 - (1 - 1/s)^s)(1 - 2\epsilon)\text{OPT}_s(F)$ (See the proof of 7).

The averaging oracle is monotone and submodular [31] when reachability functions are as in live-edge models. Unfortunately our median-of-averages oracle which facilitates tighter bounds on the number of simulations is monotone but may not be submodular even for models where the averaging oracle is submodular. Generally when this is the case, greedy may fail (as highlighted in recent work by Balkanski et al [2]).

Fortunately, greedy is effective on a function \hat{F} that is monotone but not necessarily submodular as long as \hat{F} “closely approximates” a monotone submodular F in that marginal contributions of the form

$$F(u \mid S) := F(S \cup \{u\}) - F(S)$$

are approximated well by $\hat{F}(u \mid S)$ [14]. We apply this to establish the following lemma:

29:14 Sample Complexity Bounds for Influence Maximization

► **Lemma 10.** *The greedy algorithm applied to a function \hat{F} that is monotone and provides a uniform ϵ_A -approximation of a monotone submodular function F where $\epsilon_A = \frac{\epsilon(1-\epsilon)}{14s}$ returns a set T such that $F(T) \geq (1 - (1 - 1/s)^s)(1 - \epsilon)\text{OPT}_s(F)$.*

Our proof of Lemma 10 generally applies to an approximate oracle \hat{F} of any monotone submodular function F and is presented in Appendix C. For approximate IM we obtain the following as a corollary:

► **Theorem 11.** *Consider a submodular SDM $\mathcal{G}(V, H)$ that for some $c \geq 1$ satisfies the variance bound (8). Consider a median-of-averages oracle constructed with $O(\epsilon^{-2}s^3c \ln \frac{n}{\delta})$ simulations of \mathcal{G} arranged as $r = O(s \ln \frac{n}{\delta})$ pools with $\ell = O(\epsilon^{-2}s^2c)$ simulations each. Then with probability $1 - \delta$, the set T that contains the first s nodes returned by greedy on the oracle satisfies $F(T) \geq (1 - (1 - 1/s)^s)(1 - \epsilon)\text{OPT}_s^\tau$.*

Proof. From Lemma 6, with appropriate constants, this configuration provides us with $(\epsilon/(14s), \delta)$ approximation guarantees. From Lemma 10 greedy provides the stated approximation ratio. ◀

Greedy on the median-of-averages oracle can be implemented generically for any SDM \mathcal{G} by explicitly maintaining the reachability sets $\text{Reach}(\phi_{ij}, \{v\} \cup S)$ for all nodes $v \in V$ in each simulation ϕ_{ij} as the greedy selects nodes into the seed set S . For each step, we compute the oracle value (see (9)) and select v for which the value for $\{v\} \cup S$ is maximized:

$$\arg \max_{v \in V \setminus S} \widehat{\text{mA}}^\tau(\{v\} \cup S).$$

We obtain approximation guarantees, however, only when the conditions of monotone submodular influence function and variance bounds are satisfied. For specific families of models, we can consider tailored efficient implementations that incrementally maintain reachability sets and values.

For live-edge models with additive utility (4) we consider an implementation of greedy on a median-of-averages oracle. This can be done by explicit maintenance of reachability sets or by using sketches [10, 14, 15, 11] (see Section 5.1). We obtain the following bounds (proof is deferred to Appendix Section D).

► **Theorem 12.** *Let \mathcal{G} be a live-edge model with an additive utility function (4) that satisfies the variance bound (8). Then greedy on median-of-averages oracle can be implemented with explicit reachability sets in time*

$$O(\epsilon^{-2}s^3c \ln \left(\frac{n}{\delta}\right) \bar{m}n), \quad (10)$$

where \bar{m} is the average number of edges per simulation (For an IC model, $c = \tau$ and $E[\bar{m}] = \sum_{e \in \mathcal{E}} p_e$). When using sketches, the time bound is

$$O(\epsilon^{-2}s^3 \ln \frac{n}{\delta} (c\bar{m} + s(m^* + ns) \ln n)), \quad (11)$$

where $m^* = \sum_v \max_{ij} d_v(E_{ij})$. For an IC model, $c = \tau$ and $m^* = \sum_e p_e$ in expectation.

Conclusion

We explore the “sample complexity” of IM on stochastic diffusion models and show that an approximate maximizer (within a small relative error) can be recovered from a small number of simulations as long as the variance is appropriately bounded. We establish the variance

bound for the large class of strongly submodular stochastic diffusion models. This includes IC models (where edges are drawn independently) and IGT models (where node thresholds are drawn independently) and natural extensions that allow for some dependencies. Our sample complexity bound significantly improves over the previous bounds by replacing the linear dependence in the number of nodes by a logarithmic dependence on the number of nodes and linear dependence on the length of the activation paths (which are usually very short). An interesting question for future work is to address the gap between the sample complexity and the larger number of simulations currently needed for greedy maximization.

References

- 1 N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *J. Comput. System Sci.*, 58:137–147, 1999.
- 2 E. Balkanski, A. Rubinfeld, and Y. Singer. The limitations of optimization from samples. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, 2017. doi:10.1145/3055399.3055406.
- 3 C. Borg, M. Brautbar, J. Chayes, and B. Lucier. Maximizing Social Influence in Nearly Optimal Time. In *SODA*, 2014.
- 4 W. Chen, L. V. S. Lakshmanan, and C. Castillo. *Information and Influence Propagation in Social Networks*. Morgan & Claypool, 2013.
- 5 W. Chen, W. Lu, and Y. Zhang. Time-critical influence maximization in social networks with time-delayed diffusion process. In *AAAI*, 2012.
- 6 W. Chen, C. Wang, and Y. Wang. Scalable Influence Maximization for Prevalent Viral marketing in Large-Scale Social Networks. In *KDD*. ACM, 2010.
- 7 W. Chen, Y. Wang, and S. Yang. Efficient influence maximization in social networks. In *KDD*. ACM, 2009.
- 8 W. Chen, Y. Yuan, and L. Zhang. Scalable Influence Maximization in Social Networks under the Linear Threshold Model. In *ICDM*, 2010.
- 9 Wei Chen, Tian Lin, Zihan Tan, Mingfei Zhao, and Xuren Zhou. Robust Influence Maximization. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*. ACM, 2016. doi:10.1145/2939672.2939745.
- 10 E. Cohen. Size-Estimation Framework with Applications to Transitive Closure and Reachability. *J. Comput. System Sci.*, 55:441–453, 1997.
- 11 E. Cohen. All-Distances Sketches, Revisited: HIP estimators for massive graphs analysis. *TKDE*, 2015. arXiv:1306.3284.
- 12 E. Cohen. Multi-Objective Weighted Sampling. In *HotWeb*. IEEE, 2015. full version: arXiv:1509.07445.
- 13 E. Cohen, S. Chechik, and H. Kaplan. Clustering Small Samples with Quality Guarantees: Adaptivity with One2all pps. In *AAAI*, 2018. arXiv:1706.03607.
- 14 E. Cohen, D. Delling, T. Pajor, and R. F. Werneck. Sketch-based Influence Maximization and Computation: Scaling up with Guarantees. In *CIKM*. ACM, 2014. arXiv:1408.6282.
- 15 E. Cohen, D. Delling, T. Pajor, and R. F. Werneck. Distance-Based Influence in Networks: Computation and Maximization. Technical Report cs.SI/1410.6976, arXiv, 2015. arXiv:1410.06976.
- 16 E. Cohen, N. Gossuag, and H. Kaplan. Processing Top-k Queries from Samples. In *Proceedings of the 2006 ACM conference on Emerging network experiment and technology (CoNext)*. ACM, 2006.
- 17 P. Domingos and M. Richardson. Mining the network value of customers. In *KDD*. ACM, 2001.
- 18 N. Du, Y. Liang, M-F. Balcan, and L. Song. Influence Function Learning in Information Diffusion Networks. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML'14*, 2014.

29:16 Sample Complexity Bounds for Influence Maximization

- 19 N. Du, L. Song, M. Gomez-Rodriguez, and H. Zha. Scalable Influence Estimation in Continuous-Time Diffusion Networks. In *NIPS*. NIPS, 2013.
- 20 D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, New York, NY, USA, 2010.
- 21 U. Feige. A threshold of $\ln n$ for approximating set cover. *J. Assoc. Comput. Mach.*, 45:634–652, 1998.
- 22 M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the Temporal Dynamics of Diffusion Networks. In *ICML*, 2011.
- 23 M. Gomez-Rodriguez, J. Leskovec, and A. Krause. Inferring Networks of Diffusion and Influence. In *KDD*, 2010.
- 24 A. Goyal, F. Bonchi, and L. V. S. Lakshmanan. Learning influence probabilities in social networks. In *WSDM*, 2010.
- 25 A. Goyal, W. Lu, and L.V.S. Lakshmanan. CELF++: Optimizing the Greedy Algorithm for Influence Maximization in Social Networks. In *WWW*. ACM, 2011.
- 26 M. Granovetter. Threshold Models of Collective Behavior. *American Journal of Sociology*, 83(6), 1978.
- 27 Xinran He and David Kempe. Robust Influence Maximization. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16. ACM, 2016. doi:10.1145/2939672.2939760.
- 28 K. Huang, S. Wang, G. S. Bevilacqua, X. Xiao, and L. K. S. Lakshmanan. Revisiting the Stop-and-Stare Algorithms for Influence Maximization. *PVLDB*, 10, 2017.
- 29 M. O. Jackson. *Social and economic networks*. Princeton University Press, 2010.
- 30 K. Jung, W. Heo, and W. Chen. IRIE: Scalable and robust influence maximization in social networks. In *ICDM*. ACM, 2012.
- 31 D. Kempe, J. M. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *KDD*. ACM, 2003.
- 32 J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and Glance N. Cost-effective outbreak detection in networks. In *KDD*. ACM, 2007.
- 33 B. Liu, G. Cong, X. Dong, and Y. Zeng. Time Constrained Influence Maximization in Social Networks. In *ICDM*, 2012.
- 34 B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause. Distributed Submodular Maximization: Identifying Representative Elements in Massive Data. In *NIPS*, 2013.
- 35 Elchanan Mossel and Sébastien Roch. On the submodularity of influence in social networks. In *STOC*, 2007.
- 36 Elchanan Mossel and Sebastien Roch. Submodularity of Influence in Social Networks: From Local to Global. [arXiv:math/0612046](https://arxiv.org/abs/math/0612046), July 2009. (Accessed on 08/16/2019).
- 37 G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of the approximations of maximizing submodular set functions. *Mathematical Programming*, 14, 1978.
- 38 H. T. Nguyen, M. T. Thai, and T. N. Dinh. Stop-and-Stare: Optimal Sampling Algorithms for Viral Marketing in Billion-scale Networks. In *SIGMOD*. ACM, 2016.
- 39 H. T. Nguyen, M. T. Thai, and T. N. Dinh. A Billion-Scale Approximation Algorithm for Maximizing Benefit in Viral Marketing. *IEEE/ACM Trans. Netw.*, 25(4), 2017.
- 40 M. Richardson and P. Domingos. Mining knowledge-sharing sites for viral marketing. In *KDD*. ACM, 2002.
- 41 K. Saito, R. Nakano, and M. Kimura. Prediction of Information Diffusion Probabilities for Independent Cascade Model. In *Knowledge-Based Intelligent Information and Engineering Systems (KES)*, volume 5179 of *Lecture Notes in Computer Science*. Springer, 2008.
- 42 Y. Tang, Y. Shi, and X. Xiao. Influence Maximization in Near-Linear Time: A Martingale Approach. In *SIGMOD*, 2015.
- 43 Y. Tang, X. Xiao, and Y. Shi. Influence Maximization: Near-Optimal Time Complexity Meets Practical Efficiency. In *SIGMOD*, 2014.
- 44 J. Travers and S. Milgram. An Experimental Study of the Small World Problem. *Sociometry*, 32:425–443, December 1969. doi:10.2307/2786545.

A Variance Upper Bound: Proof of Theorem 2

In this section we prove Theorem 2 which upper bounds the variance in strongly submodular SDMs. We start by bounding the variance in a more basic setting of a submodular function over a random subset in Section A.1 (Theorem 14). This will be an ingredient in our main proof provided in Section A.3.

We will be using the following basic tools:

► **Lemma 13.** *If X, Y are two random variables on the same probability space and the variance of Y is finite, then:*

$$E[Y] = E[E[Y|X]] \quad (\text{total expectation})$$

$$\text{Var}[Y] = E[\text{Var}[Y|X]] + \text{Var}[E[Y|X]] \quad (\text{total variance})$$

where $E[Y|X]$ is a random variable that gets the expectation of Y conditioned the value of X and $\text{Var}[Y|X]$ is a random variable that gets the variance of Y conditioned the value of X .

When X is a Bernoulli random variable $X \sim \text{Ber}(p)$ then Lemma 13 gives that

$$\text{Var}[Y] = E[\text{Var}[Y|X]] + \text{Var}[E[Y|X]] = pV_1 + (1-p)V_0 + p(1-p)(E_1 - E_0)^2.$$

where $V_0 = \text{Var}[Y|X=0]$, $V_1 = \text{Var}[Y|X=1]$, $E_0 = E[Y|X=0]$, and $E_1 = E[Y|X=1]$.

A.1 Submodular Monotone Functions on Random Subsets

Let $S = \{a_i\}_{1 \leq i \leq t}$ be a set with t elements and let $P = \{p_i\}_{1 \leq i \leq t}$ be a set of t probabilities such that p_i is associated with the element a_i . Let X be a random subset of S that contains a_i with probability p_i independently for each $i = 1, \dots, t$. That is

$$\forall A \in 2^S : \Pr[X = A] = \prod_{a_i \in A} p_i \prod_{a_i \notin A} (1 - p_i).$$

We say that X is a *random subset of S using probabilities P* .

A *submodular monotone function f over S* is a function with the following properties:

1. $f : 2^S \rightarrow \mathbb{R}^+$
2. For every $A, B \subset S$ with $A \subset B$ and for every $x \in S \setminus B$ we have that $f(A \cup x) - f(A) \geq f(B \cup x) - f(B)$
3. $A \subset B \Rightarrow f(A) \leq f(B)$

For any singleton $\{a\} \in S$ we write $f(a)$ instead of $f(\{a\})$. Let $M_f = \max_i f(a_i) - f(\emptyset)$. Our purpose in this subsection is to establish the following:

► **Theorem 14.** *Let X be a random subset of S using probabilities P and let f be a submodular monotone function. Then*

$$\text{Var}[f(X)] \leq M_f E[f(X) - f(\emptyset)].$$

We give the following additional definitions and lemmas before proving this theorem.

29:18 Sample Complexity Bounds for Influence Maximization

Let $S_{-i} = S \setminus \{a_i\}$ and let $P_{-i} = P \setminus \{p_i\}$. We define X_{-i} to be a random subset of S_{-i} using the probabilities P_{-i} . Let f_i^0, f_i^1 be submodular functions over S_{-i} defined by $f_i^0(A) = f(A)$ and $f_i^1(A) = f(A \cup \{a_i\})$. Let

$$E_i^1 = E[f_i^1(X_{-i})] = \sum_{A \in 2^{S \setminus \{a_i\}}} \Pr[X_{-i} = A] f_i^1(A).$$

and

$$E_i^0 = E[f_i^0(X_{-i})] = \sum_{A \in 2^{S \setminus \{a_i\}}} \Pr[X_{-i} = A] f_i^0(A).$$

By our definitions $E[f(X_{-i})] = E_i^0$ and from total expectation (Lemma 13), $E[f(X)] = p_i E_i^1 + (1 - p_i) E_i^0$.

► **Lemma 15.** *let f be a submodular monotone function over S and X a random subset of S using probabilities P . Then,*

$$\forall i : E_i^1 - E_i^0 \leq f(a_i) - f(\emptyset) \leq M_f.$$

Proof. Since X is obtained by drawing the elements in S independently it follows that

$$\begin{aligned} E_i^1 - E_i^0 &= \sum_{A \in 2^{S-i}} \Pr[X_{-i} = A] [f(A \cup \{a_i\}) - f(A)] \leq \underbrace{\phantom{\sum_{A \in 2^{S-i}} \Pr[X_{-i} = A] [f(A \cup \{a_i\}) - f(A)]}}_{\text{submodularity}} \\ &\leq \sum_{A \in 2^{S \setminus a_i}} \Pr[X_{-i} = A] [f(a_i) - f(\emptyset)] \leq f(a_i) - f(\emptyset) \leq M_f. \quad \blacktriangleleft \end{aligned}$$

► **Lemma 16.** *for any submodular monotone function f over S and for any index i we have that $M_{f_i^0} \leq M_f$ and $M_{f_i^1} \leq M_f$.*

Proof. The first inequality follows immediately from our definition since

$$M_{f_i^0} = \max_{j \neq i} f_i^0(a_j) - f_i^0(\emptyset) = \max_{j \neq i} f(a_j) - f(\emptyset) \leq M_f.$$

For the second inequality we use submodularity as follows

$$M_{f_i^1} = \max_{j \neq i} f_i^1(a_j) - f_i^1(\emptyset) = \max_{j \neq i} f(\{a_j \cup a_i\}) - f(a_i) \leq \underbrace{\phantom{\max_{j \neq i} f(\{a_j \cup a_i\}) - f(a_i)}}_{\text{sub-modularity}} \max_j f(a_j) - f(\emptyset) = M_f. \quad \blacktriangleleft$$

We are now ready for the proof of Theorem 14.

Proof of Theorem 14. The proof is by induction on the size of S .

Base case. Let $S = \{a_1\}, P = \{p_1\}$ we have that

$$E[f(X)] = p_1 f(a_1) + (1 - p_1) f(\emptyset),$$

and

$$\begin{aligned} \text{Var}[f(X)] &= E[f(X)^2] - E^2[f(X)] = p_1 f^2(a_1) + (1 - p_1) f^2(\emptyset) - [p_1 f(a_1) + (1 - p_1) f(\emptyset)]^2 \\ &= p_1(1 - p_1) [f(a_1) - f(\emptyset)]^2 \leq M_f p_1(1 - p_1) [f(a_1) - f(\emptyset)]. \end{aligned}$$

It is left to prove that $\mathbb{E}[f(X) - f(\emptyset)] \geq p_1(1 - p_1)[f(a_1) - f(\emptyset)]$, and indeed we have that

$$\mathbb{E}[f(X) - f(\emptyset)] = p_1[f(a_1) - f(\emptyset)] \geq p_1(1 - p_1)[f(a_1) - f(\emptyset)].$$

Inductive Step. Assume the lemma holds for sets of size ℓ and any submodular function f and probabilities P . For a set S with $\ell + 1$ elements and a submodular function f over S . Let $j \leq i$ be an arbitrary index.

From the total variance formula in Lemma 13 we know that

$$\text{Var}[f(X)] = p_j V_j^1 + (1 - p_j) V_j^0 + p_j(1 - p_j) [E_j^1 - E_j^0]^2, \quad (12)$$

where $E_j^1 = \mathbb{E}[f_j^1(X_{-j})]$, $E_j^0 = \mathbb{E}[f_j^0(X_{-j})]$, $V_j^1 = \text{Var}[f_j^1(X_{-j})]$, and $V_j^0 = \text{Var}[f_j^0(X_{-j})]$. By applying the induction hypothesis to S_{-j} with probabilities P_{-j} and $|S_{-j}| = \ell$ and f_j^0 and f_j^1 we get that

$$V_j^0 \leq \underbrace{\mathbf{M}_{f_j^0}[E_j^0 - f_j^0(\emptyset)]}_{\text{Lemma 16}} \leq \mathbf{M}_f[E_j^0 - f(\emptyset)],$$

and

$$V_j^1 \leq \underbrace{\mathbf{M}_{f_j^1}[E_j^1 - f_j^1(\emptyset)]}_{\text{Lemma 16}} \leq \mathbf{M}_f[E_j^1 - f(a_j)].$$

Substituting these bounds in Equation (12) we get that

$$\begin{aligned} \text{Var}[f(X)] &\leq p_j \mathbf{M}_f[E_j^1 - f(a_j)] + (1 - p_j) \mathbf{M}_f[E_j^0 - f(\emptyset)] + p_j(1 - p_j) [E_j^1 - E_j^0]^2 \\ &= \mathbf{M}_f[p_j E_j^1 + (1 - p_j) E_j^0 - f(\emptyset)] + p_j(1 - p_j) [E_j^1 - E_j^0]^2 - p_j \mathbf{M}_f[f(a_j) - f(\emptyset)] \\ &\quad \underbrace{=}_{\text{total expectation}} \mathbf{M}_f[\mathbb{E}[f(X)] - f(\emptyset)] + p_j(1 - p_j) [E_j^1 - E_j^0]^2 - p_j \mathbf{M}_f[f(a_j) - f(\emptyset)] \\ &\quad \underbrace{\leq}_{\text{Lemma 15}} \mathbf{M}_f[\mathbb{E}[f(X)] - f(\emptyset)] + p_j \mathbf{M}_f[f(a_j) - f(\emptyset)] [(1 - p_j) - 1] \\ &\leq \mathbf{M}_f[\mathbb{E}[f(X)] - f(\emptyset)]. \end{aligned} \quad \blacktriangleleft$$

A.2 Properties of Reduced Diffusion Models

We establish some properties of reduced independent SDMs that are needed for our upper bound.

We first show that influence values of nodes in a reduced model can only be lower than respective values in the original model:

► **Lemma 17.** *Let $\mathcal{G}'(V \setminus T, H')$ be a reduction of a model $\mathcal{G}(V, H)$.*

$$\text{For all } v \in V \setminus T \text{ and } t \geq 0, \mathcal{I}_{\mathcal{G}'}^t(v) \leq \mathcal{I}_{\mathcal{G}}^t(v). \quad (13)$$

Proof. Note that \mathcal{G}' is obtained from \mathcal{G} by removing nodes. Therefore respective reachability sets given ϕ are such that those in \mathcal{G}' can only be subsets of those in \mathcal{G} :

$$\text{Reach}_{\mathcal{G}'}(v, \phi') \subset \text{Reach}_{\mathcal{G}}(v, \phi).$$

Then from monotonicity and submodularity of H we get

$$H'(\text{Reach}_{\mathcal{G}'}(v, \phi')) \leq H(\text{Reach}_{\mathcal{G}}(v, \phi)) \leq H(\text{Reach}_{\mathcal{G}}(v, \phi)).$$

29:20 Sample Complexity Bounds for Influence Maximization

(second inequality follows from monotonicity and submodularity of H so that for all $A \subset V \setminus T$, $H'(A) \leq H(A)$.) Therefore,

$$I_{\mathcal{G}'}^t(v) = \mathbb{E}[\text{VReach}_{\mathcal{G}'}^t(v)] \leq \mathbb{E}[\text{VReach}_{\mathcal{G}}^t(v)] = I_{\mathcal{G}}^t(v). \quad \blacktriangleleft$$

A convenient property is that reduction preserves strong monotone submodularity:

► **Lemma 18.** *A reduction of a strongly monotone submodular model is also strongly monotone submodular.*

Proof. A reduced model with respect to T_2 of a reduced model of \mathcal{G} with respect to T_1 is a reduced model of \mathcal{G} with respect to $T_1 \cup T_2$. Also note that the reduced utility function H' is also monotone and submodular. ◀

We next show that IC or IGT models with submodular utility are closed under reduction:

► **Theorem 19.** *IC and IGT models with submodular utility are strongly submodular SDMs.*

Proof. We first show that IC/IGT models are independent SDMs. In the introduction we expressed IC and IGT models as SDMs: A live-edge model is expressed as an SDM using $\phi_v(T) = 1$ if and only if there is an edge from a node in T to v . The model is independent if for all v the edges incoming to v are independent of all other edges. In IC models all edges are independent and hence IC models are independent SDMs. Recall (from the Introduction) that an IGT model is expressed as an SDM using $\phi_v(T) := \text{Indicator}(\theta_v \leq f_v(T))$. In an IGT model the thresholds θ_v are independent random variables, and hence ϕ_v are independent. Hence, an IGT model is an independent SDM. Submodularity of influence when utility is submodular is established for IC models in [31] and for IGT models in [35].

Reduction of any model preserves submodularity of the utility and in particular this holds for reduced IC/IGT models. What remains to show is that a reduced IC/IGT model is also an IC/IGT model (respectively). This would conclude the proof of strong submodularity since any IC and IGT models with submodular utility has a submodular influence functions.

To establish this remaining claim we consider IC/IGT models and express the reduction in terms of the activation functions as one in terms of the respective family of models.

We first consider IC models. The reduced IC model $\mathcal{G}'(V \setminus T, \mathcal{E} \setminus (V \times T \cup T \times V))$ is obtained from $\mathcal{G}(V)$ by deleting the nodes T and their incident edges and keeping p_e on remaining edges. This is clearly an IC model. It remains to show that this is equivalent to the reduction of the distribution of activation functions. The conditioning that $\phi_v(T) = 0$ is equivalent to live-edge set E with no edges from T to V . For such edge set for any $S \subset V \setminus (T \cup \{v\})$ we have $\phi'_v(S) = \phi_v(S \cup T) = \phi_v(S)$ which corresponds to E having at least one edge from S to v . From independence of edges, the conditional distribution is also independent and retains the same inclusion probabilities.

We next consider IGT models. The reduction $\mathcal{G}'(V \setminus T, \{f'_v\})$ in terms of activation functions distribution is equivalent to functions is equivalent to modifying the functions so that

$$f'_v(S) := f_v(S \cup T) - f_v(T).$$

The reduced model is clearly an IGT model. The conditioning that $\phi_v(T) = 0$ means that $\theta_v > f_v(T)$. Therefore, the conditional distribution of θ_v provided it was not activated in the first step is uniform on $[f_v(T), 1]$. The probability that $\theta_v > f_v(S \cup T)$ given this conditioning is equal to the probability that $\theta'_v > f_v(S \cup T) - f_v(T) = f'_v(S)$. ◀

A.3 Upper Bound on the Variance in Strongly Submodular SDM

Let $\mathcal{G}(V, H)$ be a τ -stepped diffusion model. We denote by $M_{\mathcal{G}}^{\tau}(\bar{T})$ the maximum influence of a single node in \mathcal{G} that is not included in T :

$$M_{\mathcal{G}}^{\tau}(\bar{T}) = \max_{v \in V \setminus T} I_{\mathcal{G}}^{\tau}(v)$$

As before, we omit \mathcal{G} if it can be understood from the context. We prove the following theorem which is a restatement of Theorem 2.

► **Theorem 20.** *Let $\mathcal{G}(V, H)$ be a strongly submodular SDM. Then for any $\tau \geq 0$ and a set of nodes $T \subset V$:*

$$\text{Var}[\mathbf{R}^{\tau}(T)] \leq \tau M^{\tau-1}(\bar{T}) I^{\tau}(T).$$

The remaining part of this Subsection contains the proof of the Theorem.

Let T be a set of nodes, and let

$$N(T) = \{v \in V \setminus T \mid \Pr[\phi_v(T) = 1] > 0\}$$

be the nodes that have nonzero probability to be activated if T is active. For the special case of IC models, $N(T) = \{v \notin T \mid \exists(u, v) \in \mathcal{E}, u \in T\}$ is the set of outgoing neighbors of T .

We first consider the case where $N(T)$ is empty. In this case, $\text{Reach}^{\tau}(T) = T$ for all $\tau \geq 0$. Therefore, $\text{Var}[\mathbf{R}^{\tau}(T)] = 0$, $I^{\tau}(T) = H(T) \geq 0$, and $M^{\tau-1}(\bar{T}) = 0$ and the claim holds.

We now assume that $N(T)$ is not empty and give a proof by induction on τ .

A.3.1 Base case ($\tau = 1$)

Let

$$p_v := \Pr[\phi_v(T) = 1]$$

be the probability that node v is activated in step 1 provided that the set of nodes T was active at step 0. From independence of the model, the events of activating different nodes $v \in N(T)$ at step 1 are independent. We have that the set of nodes that is active at step 1 is a random subset S of $N(T)$ with probabilities $\{p_v\}$ as defined in Subsection A.1. Moreover, from monotonicity and submodularity of H , the function $f(S) := H(T \cup S) - H(T)$ is monotone and submodular with $f(\emptyset) = 0$. We can therefore apply Theorem 14 to bound the variance of $f(S)$:

$$\text{Var}[f(S)] \leq M_f \mathbf{E}[f(S)]. \tag{14}$$

We now note that

$$\mathbf{E}[f(S)] = \mathbf{E}[H(T \cup S)] - H(T) = I^1(T) - I^0(T)$$

and

$$\text{Var}[f(S)] = \text{Var}[H(T \cup S)] = \text{Var}[\mathbf{R}^1(T)].$$

For all $v \in N(T)$ we have $f(v) = H(T \cup \{v\}) - H(T) \leq H(v) = I^0(v)$. Therefore

$$M_f := \max_{v \in N(T)} f(v) \leq \max_{v \in N(T)} I^0(v) = M^0(\bar{T}).$$

Substituting in (14) we obtain the claim

$$\text{Var}[\mathbf{R}^1(T)] \leq M^0(\bar{T}) I^1(T).$$

A.3.2 Inductive step

We define $\text{Reach}_{\mathcal{G}}^t(T | A)$ to be the random variable that is the t -steps reachability of T in a diffusion on \mathcal{G} seeded with T that is conditioned on the event that exactly the nodes $A \subset N(T)$ (and no other nodes) are activated in step 1. Equivalently, we condition on ϕ such that for $v \setminus (T \cup A)$, $\phi_v(T) = 0$ and for $v \in A$, $\phi_v(T) = 1$. We respectively define $\mathbb{R}_{\mathcal{G}}^t(T | A)$ to be the random variable $H(\text{Reach}_{\mathcal{G}}^t(T | A))$. From definition, we have

$$\mathbb{I}_{\mathcal{G}}^t(T) = \sum_{A \subset N(T)} \Pr[\text{Reach}_{\mathcal{G}}^1(T) = A \cup T] \mathbb{E}[\mathbb{R}_{\mathcal{G}}^t(T | A)] = \mathbb{E}_A \mathbb{E}[\mathbb{R}_{\mathcal{G}}^t(T | A)]. \quad (15)$$

We consider the reduced model \mathcal{G}' of \mathcal{G} with respect to T and show that the conditioned $t \geq 1$ steps diffusion from T in \mathcal{G} is equivalent to the unconditioned $t - 1$ steps diffusion from A in \mathcal{G}' :

► **Lemma 21.** *For any $A \subset N(T)$ and $t \geq 1$, the random variables $\text{Reach}_{\mathcal{G}'}^{t-1}(A)$ and $\text{Reach}_{\mathcal{G}}^t(T|A) \setminus \{T\}$ have identical distribution over subsets. The random variables $\mathbb{R}_{\mathcal{G}'}^{t-1}(A)$ and $\mathbb{R}_{\mathcal{G}}^t(T|A) - H(T)$ have identical distributions over values.*

Proof. We first consider $t = 1$. For a draw of conditioned activation functions we have $\text{Reach}_{\mathcal{G}}^1(T|A) = T \cup A$. By definition, we also have $\text{Reach}_{\mathcal{G}'}^0(A) = A$ and the claim holds.

We next consider $t > 1$. We first observe that in both situations, (i) the reduced model \mathcal{G}' when seeded with A and (ii) the conditioned diffusion in \mathcal{G} seeded with T such that the nodes A are activated in the first step, the progression is determined only by the activation functions on the nodes $V \setminus (T \cup A)$.

We next argue that the distribution of activation functions projected on the nodes $V \setminus (T \cup A)$ is the same in both situations. From independence of \mathcal{G} it suffices to consider separately the activation functions of each node. From definition of a reduced model, we draw for each $v \in V \setminus T$, $\phi_v \sim \mathcal{G}$ conditioned on $\phi_v(T) = 0$. This is exactly what we get for the conditioned diffusion in \mathcal{G} .

We can thus match the supports (sets of activations functions) in both situations so that ϕ and ϕ' are matched when the projections on $V \setminus (T \cup A)$ is the same. The starting points are at steps 0 of the reduced model and step 1 of the conditioned process is A , the progression of new activations is thus the same. Therefore, for any step $t \geq 1$,

$$\text{Reach}_{\mathcal{G}}^t(T | A, \phi) \setminus T = \text{Reach}_{\mathcal{G}'}^{t-1}(A, \phi')$$

and the first claim follows.

For the second claim, note that $\mathbb{R}_{\mathcal{G}}^t(T|A) = H(\text{Reach}_{\mathcal{G}}^t(T|A))$ and thus

$$\mathbb{R}_{\mathcal{G}'}^{t-1}(A) = H'(\text{Reach}_{\mathcal{G}'}^{t-1}(A)) = H(\text{Reach}_{\mathcal{G}'}^{t-1}(A) \cup T) - H(T) = H(\text{Reach}_{\mathcal{G}}^t(T|A)) - H(T)$$

where the equalities are those of distributions. ◀

As immediate corollaries we can relate expectations and variance of as follows:

$$\begin{aligned} \mathbb{I}_{\mathcal{G}'}^{\tau-1}(A) &= \mathbb{E}[\mathbb{R}_{\mathcal{G}'}^{\tau-1}(A)] = \mathbb{E}[H'(\text{Reach}_{\mathcal{G}'}^{\tau-1}(A))] = \mathbb{E}[H(\text{Reach}_{\mathcal{G}}^{\tau}(T|A))] - H(T) \quad (16) \\ &= \mathbb{E}[\mathbb{R}_{\mathcal{G}}^{\tau-1}(T|A)] - H(T). \end{aligned}$$

$$\begin{aligned} \text{Var}[\mathbb{R}_{\mathcal{G}'}^{\tau-1}(A)] &= \text{Var}[H'(\text{Reach}_{\mathcal{G}'}^{\tau-1}(A))] = \text{Var}[H'(\text{Reach}_{\mathcal{G}}^{\tau}(T|A) \setminus \{T\})] \quad (17) \\ &= \text{Var}[H(\text{Reach}_{\mathcal{G}}^{\tau}(T|A)) - H(T)] = \text{Var}[H(\text{Reach}_{\mathcal{G}}^{\tau}(T|A))] = \text{Var}[\mathbb{R}_{\mathcal{G}}^{\tau}(T|A)]. \end{aligned}$$

A.3.2.1 Total Variance

We define the random variable A to be the subset of $N(T)$ which is activated after the first step. Note that A is a random subset of $N(T)$ using probabilities p_v for $v \in N(T)$ as defined in Section A.1. By the total variance formula we get that

$$\text{Var}[\mathbb{R}_{\mathcal{G}}^{\tau}(T)] = \text{Var}_A[\mathbb{E}[\mathbb{R}_{\mathcal{G}}^{\tau}(T|A)]] + \mathbb{E}_A[\text{Var}[\mathbb{R}_{\mathcal{G}}^{\tau}(T|A)]] . \quad (18)$$

We bound the total variance by separately bounding the two terms.

A.3.2.2 Bound on the first term of the total variance

We consider the reduced model \mathcal{G}' with respect to T and a restriction of the influence function $\mathbb{I}_{\mathcal{G}'}^{\tau-1}$ to the domain that is subsets $A \subseteq N(T)$:

$$f(A) := \mathbb{I}_{\mathcal{G}'}^{\tau-1}(A).$$

From Lemma 21, this function represents the expected marginal utility value of nodes which are not in T that are activated after τ steps if we activate T at step 0 and the set A at step 1.

We first observe that f is monotone and submodular. This because strong monotone submodularity of our model implies that the reduced model is also strongly monotone and submodular, and a restriction of a monotone and submodular function is also monotone and submodular. We establish two helpful properties of f . First,

$$f(\emptyset) = 0 , \quad (19)$$

which holds for any influence function. Second, using Lemma 17 we obtain

$$\max_{v \in N(T)} f(v) \stackrel{(13)}{\leq} \max_{v \in N(T)} \mathbb{I}_{\mathcal{G}}^{\tau-1}(v) \leq \mathbb{M}^{\tau-1}(\bar{T}) . \quad (20)$$

We are now ready to bound the first term of the total variance (18). Our monotone submodular function f and the random subset A using probabilities p_v satisfy the conditions of Theorem 14.

$$\begin{aligned} \text{Var}_A[\mathbb{E}[\mathbb{R}^{\tau}(T|A)]] &= \text{Var}_A[f(A) + H(T)] = \text{Var}_A[f(A)] \stackrel{\text{Theorem 14}}{\leq} \left(\max_{v \in N(T)} f(v) \right) \mathbb{E}_A [f(A) - f(\emptyset)] \\ &\stackrel{(19), (20)}{\leq} \mathbb{M}^{\tau-1}(\bar{T}) \mathbb{E}_A [\mathbb{E}[\mathbb{R}^{\tau}(T|A) - H(T)]] = \mathbb{M}^{\tau-1}(\bar{T}) (\mathbb{I}_{\mathcal{G}}^{\tau}(T) - H(T)) \end{aligned} \quad (21)$$

A.3.2.3 Bound on the second term of the total variance

We next bound the second term of (18) which is the expectation of the variance conditioned on A :

29:24 Sample Complexity Bounds for Influence Maximization

$$\begin{aligned}
\mathbb{E}_A[\text{Var}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | A)]] &= \sum_{S \subset N(T)} \Pr[A = S] \text{Var}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | S)] \\
&= \sum_{S \subset N(T)} \Pr[A = S] \mathbb{E}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | S) - H(T)] \frac{\text{Var}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | S)]}{\mathbb{E}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | S) - H(T)]} \\
&\leq \max_{S \subset N(T)} \frac{\text{Var}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | S)]}{\mathbb{E}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | S) - H(T)]} \sum_{S \subset N(T)} \Pr[A = S] \mathbb{E}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | S) - H(T)] \\
&= \mathbb{E}_A [\mathbb{E}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | A) - H(T)]] \max_{S \subset N(T)} \frac{\text{Var}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | S)]}{\mathbb{E}[\mathbb{R}_{\mathcal{G}}^{\tau}(T | S) - H(T)]} \\
&\stackrel{(15),(16),(17)}{=} \underbrace{(\mathbb{I}_{\mathcal{G}}^{\tau}(T) - H(T))}_{(15),(16),(17)} \max_{S \subset N(T)} \frac{\text{Var}[\mathbb{R}_{\mathcal{G}'}^{\tau-1}(S)]}{\mathbb{I}[\mathbb{R}_{\mathcal{G}'}^{\tau-1}(S)]} \\
&= (\mathbb{I}_{\mathcal{G}}^{\tau}(T) - H(T)) \frac{\text{Var}[\mathbb{R}_{\mathcal{G}'}^{\tau-1}(S')]}{\mathbb{I}[\mathbb{R}_{\mathcal{G}'}^{\tau-1}(S')]} \tag{22}
\end{aligned}$$

Where we take

$$S' = \arg \max_{S \subset N(T)} \frac{\text{Var}[\mathbb{R}_{\mathcal{G}'}^{\tau-1}(S)]}{\mathbb{I}[\mathbb{R}_{\mathcal{G}'}^{\tau-1}(S)]}$$

to be the subset that maximizes the ratio.

Using the induction hypothesis on $(\tau - 1)$ -stepped influence we get

$$\text{Var}[\mathbb{R}_{\mathcal{G}'}^{\tau-1}(S')] \leq (\tau - 1) \mathbb{M}_{\mathcal{G}'}^{\tau-2}(\bar{S}') \mathbb{I}[\mathbb{R}_{\mathcal{G}'}^{\tau-1}(S')]. \tag{23}$$

We now relate the maximum influence of nodes in the original and reduced models:

$$\mathbb{M}_{\mathcal{G}'}^{\tau-2}(\bar{S}') = \max_{v \in V \setminus (T \cup S')} \mathbb{I}_{\mathcal{G}'}^{\tau-2}(v) \stackrel{(13)}{\leq} \max_{v \in V \setminus (T \cup S')} \mathbb{I}_{\mathcal{G}}^{\tau-2}(v) \leq \mathbb{M}_{\mathcal{G}}^{\tau-2}(\overline{T \cup S'}) \leq \mathbb{M}_{\mathcal{G}}^{\tau-1}(\bar{T}). \tag{24}$$

From (22) using (23) and (24) we obtain

$$\mathbb{E}_A[\text{Var}[\mathbb{R}^{\tau}(T|A)]] \stackrel{(23)}{\leq} \underbrace{(\mathbb{I}_{\mathcal{G}}^{\tau}(T) - H(T))}_{(23)} (\tau - 1) \mathbb{M}_{\mathcal{G}'}^{\tau-2}(\bar{S}') \stackrel{(24)}{\leq} \underbrace{(\mathbb{I}_{\mathcal{G}}^{\tau}(T) - H(T))}_{(24)} (\tau - 1) \mathbb{M}_{\mathcal{G}}^{\tau-1}(\bar{T}). \tag{25}$$

A.3.2.4 Combining the bounds of the first and second terms

The claim of the Theorem follows using total variance (18) and the bounds on the first term (21) and second term (25).

B Variance lower bound construction

► **Lemma 22.** *Let \mathcal{G} be complete binary tree where each edge has probability $\frac{1}{2}$ and let $h(u)$ be the height of the node u . Then, $\mathbb{I}^{\tau}(u) = h(u)$ and $\text{Var}[\mathbb{R}^{\tau}(u)] = \frac{1}{2} \sum_{i=0}^{h(u)-1} i^2 = \frac{(h(u)-1)h(u)(2h(u)-1)}{12}$.*

Proof. By induction on the height of the of the node.

Base step ($h(u) = 1$). It is clear that $\mathbb{I}^1(u) = 1$ and $\text{Var}[\mathbb{R}^1(u)] = 0$ since u is a leaf.

Inductive step. u has two neighbors and each is reached with probability $\frac{1}{2}$. Let v_1 and v_2 be the neighbors of u and let X_1, X_2 be random variables that indicate if $(u, v_1), (u, v_2)$ were activated respectively. The variables X_1, X_2 are Bernoulli random variables with $p = \frac{1}{2}$, hence, $\mathbb{E}[X_1] = \mathbb{E}[X_2] = \frac{1}{2}$ and $\text{Var}[X_1] = \text{Var}[X_2] = \frac{1}{4}$. Since the graph is a tree, the reachabilities of v_1 and v_2 are independent random variables, so we can simply write:

$$\mathbb{R}^\tau(u) = 1 + X_1 \mathbb{R}^{\tau-1}(v_1) + X_2 \mathbb{R}^{\tau-1}(v_2)$$

The variable $\mathbb{R}^{\tau-1}(v_1)$ and $\mathbb{R}^{\tau-1}(v_2)$ are identical and $X_1, \mathbb{R}^{\tau-1}(v_1)$ and $X_2, \mathbb{R}^{\tau-1}(v_2)$ are independent random variables, Thus,

$$\begin{aligned} \mathbb{I}^\tau(u) &= \mathbb{E}[\mathbb{R}^\tau(u)] = 1 + \mathbb{E}[X_1] \mathbb{E}[\mathbb{R}^{\tau-1}(v_1)] + \mathbb{E}[X_2] \mathbb{E}[\mathbb{R}^{\tau-1}(v_2)] = 1 + 2 \frac{1}{2} \mathbb{E}[\mathbb{R}^{\tau-1}(v_1)] \\ &\stackrel{\text{induction's hypothesis}}{=} 1 + h(u) - 1 = h(u) \end{aligned}$$

The computation of the variance is similar:

$$\text{Var}[\mathbb{R}^\tau(u)] = \text{Var}[X_1 \mathbb{R}^{\tau-1}(v_1)] + \text{Var}[X_2 \mathbb{R}^{\tau-1}(v_2)] = 2 \text{Var}[X_1 \mathbb{R}^{\tau-1}(v_1)]$$

For two independent random variables A, B holds that: $\text{Var}[AB] = \text{Var}[A] \text{Var}[B] + \text{Var}[A] \mathbb{E}^2[B] + \text{Var}[B] \mathbb{E}^2[A]$, we have that:

$$\begin{aligned} \text{Var}[\mathbb{R}^\tau(u)] &= 2 (\text{Var}[X_1] \text{Var}[\mathbb{R}^{\tau-1}(v_1)] + \text{Var}[X_1] \mathbb{E}^2[\mathbb{R}^{\tau-1}(v_1)] + \text{Var}[\mathbb{R}^{\tau-1}(v_1)] \mathbb{E}^2[X_1]) \\ &= 2 \left(\frac{1}{2} \text{Var}[\mathbb{R}^{\tau-1}(v_1)] + \frac{1}{4} \mathbb{E}^2[\mathbb{R}^{\tau-1}(v_1)] \right) = \text{Var}[\mathbb{R}^{\tau-1}(v_1)] + \frac{1}{2} (h(u) - 1)^2 \\ &\stackrel{\text{induction's hypothesis}}{=} \frac{1}{2} \sum_{i=0}^{h(u)-1} i^2 = \frac{(h(u) - 1)h(u)(2h(u) - 1)}{12}. \end{aligned} \quad \blacktriangleleft$$

► **Theorem 23.** *There is a model \mathcal{G} and a set of nodes T such that $\frac{\text{Var}[\mathbb{R}^\tau(T)]}{M^\tau(T) \mathbb{I}^\tau(T)} \geq \frac{\tau}{12}$.*

Proof. Lemma 22 shows that for every node $u \in \mathcal{G}$, $\mathbb{I}^\tau(u) = h(u)$ and $\text{Var}[\mathbb{R}^\tau(u)] = \frac{(h(u)-1)h(u)(2h(u)-1)}{12}$. It follows that the root r has the largest influence $\mathbb{I}^\tau(r) = \tau$ and $\text{Var}[\mathbb{R}^\tau(u)] = \frac{(\tau-1)\tau(2\tau-1)}{12}$, Furthermore $M^\tau(\bar{r}) = \tau - 1$ since the nodes of the largest influence in $V \setminus r$ are the children of r . We conclude that:

$$\frac{\text{Var}[\mathbb{R}^\tau(r)]}{M^\tau(\bar{r}) \mathbb{I}^\tau(r)} = \frac{(\tau-1)\tau(2\tau-1)}{\tau(\tau-1)12} = \frac{2\tau-1}{12} \stackrel{\tau \geq 1}{\geq} \frac{\tau}{12}. \quad \blacktriangleleft$$

C Greedy Optimization with an Approximately Submodular Oracle

In this section we present the proof of Lemma 10. We show that our approximation guarantees imply that the application of greedy on \hat{F} generates a sequence that is an approximate greedy sequence (in the sense of Lemma 24) with respect to F .

We first state a helpful Lemma [14] that establishes that it suffices that $\hat{F}(u | S)$ to approximate the marginal contributions

$$F(u | S) := F(S \cup \{u\}) - F(S)$$

► **Lemma 24** ([14]). *Given a monotone submodular function F , an approximate greedy algorithm that for some $\epsilon \in [0, 1)$ selects at each step an element u such that $F(u | S) \geq (1 - \epsilon) \max_v F(v | S)$ has approximation ratio $\geq (1 - (1 - 1/s)^s)(1 - \epsilon)$.*

Proof. It is easy to see that the approximation ratio of ϵ -approximate greedy is $1 - (1 - (1 - \epsilon)/s)^s$. It therefore suffices to establish that this expression is larger than $(1 - (1 - 1/s)^s)(1 - \epsilon)$ for $\epsilon \in [0, 1]$. Equivalently, we need to show that for all $s \geq 2$ and $x \in [0, 1]$

$$(1 - (1 - x)/s)^s - (1 - x)(1 - 1/s)^s - x \leq 0 .$$

This follows from equality holding for $x = 0$ and $x = 1$ and the function being concave up (second derivative is positive). ◀

Proof of Lemma 10. Consider a monotone non-negative \hat{F} that is a uniform ϵ_A -approximation of a monotone non-negative F with $\epsilon_A = \frac{\epsilon(1-\epsilon)}{14s}$. By definition of ϵ -approximation (see Section 3.5), $|\hat{F}(T) - F(T)| \geq \epsilon_A \max\{F(T), \text{OPT}_1(F)\}$ for all S with $|S| \leq s$. Therefore,

$$\text{if } F(S) \geq (1 - \epsilon)\text{OPT}_1(F) \text{ then } \frac{|\hat{F}(S) - F(S)|}{F(S)} \leq \frac{\epsilon}{14s} \quad (26)$$

$$\text{if } F(S) \leq (1 - \epsilon)\text{OPT}_1(F) \text{ then } \hat{F}(S) \leq \left(1 - \frac{\epsilon}{2}\right) \text{OPT}_1(F) . \quad (27)$$

Inequality (26) follows immediately when $F(S) \geq \text{OPT}_1(F)$ because the relative error is at most $\epsilon_A \leq \frac{\epsilon}{14s}$. For $(1 - \epsilon)\text{OPT}_1(F) \leq F(S) < \text{OPT}_1(F)$ we have absolute error being at most $\epsilon_A \text{OPT}_1(S)$ which is a relative error of at most $\epsilon_A/(1 - \epsilon) \leq \frac{\epsilon}{14s}$. Inequality (27) follows from the absolute error being at most $\epsilon_A \text{OPT}_1(F)$ and $\epsilon_A \leq \epsilon/2$.

We establish that these conditions imply that greedy on \hat{F} on the prefix of the greedy sequence where $F(S) \leq \frac{3}{4}\text{OPT}_s(F)$ is actually approximate greedy (as in the conditions of Lemma 24) with respect to F . Note that $1 - (1 - 1/s)^s \geq 3/4$ for $s \geq 2$ and thus the prefix restriction does not limit generality. The claim will then follow from Lemma 24.

For $s = 1$, it follows from Equations (26) and (27), that the first element of a greedy sequence with respect to \hat{F} , $\arg \max_u \hat{F}(u)$, satisfies $\hat{F}(u) \geq (1 - \frac{\epsilon}{14})\text{OPT}_1(F)$. Therefore from the second iteration and on, we have a set S for which the relative error bounds in Equation (26) applies.

We consider the marginal contributions $F(u | S)$ for any node u . We have

$$\begin{aligned} |\hat{F}(S) - F(S)| &\leq \frac{\epsilon}{14s} F(S) \\ |\hat{F}(S \cup \{u\}) - F(S \cup \{u\})| &\leq \frac{\epsilon}{14s} F(S \cup \{u\}) = \frac{\epsilon}{14s} (F(S) + F(u | S)) . \end{aligned}$$

We use these inequalities to bound the absolute error of (any) marginal influence estimate by

$$\begin{aligned} \left| \hat{F}(u | S) - F(u | S) \right| &= \left| \hat{F}(S \cup \{u\}) - \hat{F}(S) - F(S \cup \{u\}) + F(S) \right| \quad (28) \\ &\leq \left| \hat{F}(S \cup \{u\}) - F(S \cup \{u\}) \right| + \left| \hat{F}(S) - F(S) \right| \\ &\leq \frac{\epsilon}{7s} F(S) + \frac{\epsilon}{14s} F(u | S) . \end{aligned}$$

We now consider the node $v = \arg \max_{u \in V} F(u | S)$ with maximum marginal contribution to S with respect to F and its contribution value

$$\Delta := F(v | S) \geq \frac{1}{s} (\text{OPT}_s - F(s)) .$$

Thus, when $F(S) \leq \frac{3}{4}\text{OPT}_s$,

$$\Delta \geq \frac{1}{3s}F(S). \quad (29)$$

By applying (28) to v we get that

$$\hat{F}(v | S) \geq \Delta - \frac{\epsilon}{7s}F(S) - \frac{\epsilon}{14s}\Delta$$

Therefore the node $v' = \arg \max_v \hat{F}(v' | S)$ with maximum marginal contribution according to \hat{F} satisfies

$$\hat{F}(v' | S) \geq \hat{F}(v | S) \geq \Delta - \frac{\epsilon}{7s}F(S) - \frac{\epsilon}{14s}\Delta.$$

By using (28) again, substituting (29), and using that fact that $s \geq 2$:

$$F(v' | S) \geq \Delta - 2\frac{\epsilon}{7s}F(S) - 2\frac{\epsilon}{14s}\Delta \quad (30)$$

$$\geq \Delta - \epsilon\Delta\left(\frac{1}{7s} + \frac{6}{7}\right) \geq \Delta(1 - \epsilon). \quad (31)$$

Therefore, the greedy sequence according to \hat{F} is an approximate greedy sequence according to F and satisfies the conditions of Lemma 24. Therefore the resulting sequence yields an approximation ratio at least $(1 - (1 - 1/s)^s)(1 - \epsilon)$. ◀

D Greedy for Live-Edge Models

Proof of Theorem 12. For the first bound, we explicitly maintain for each node $u \in V$, for each pool, the reachability set of u in the simulations of the pool (and its cardinality). The dominant term in the cost of computation is performing a BFS from each node in each of the $r\ell$ simulations that is truncated at distance τ . The total computation time is

$$O(r\ell\bar{m}n) = O(\epsilon^{-2}s^3c \ln\left(\frac{n}{\delta}\right)\bar{m}n), \quad (32)$$

where

$$\bar{m} = \frac{1}{\ell r} \sum_{i=1}^r \sum_{j=1}^{\ell} |E_{ij}|$$

is the average number of edges per simulation. For an IC model, $E[\bar{m}] = \sum_{e \in \mathcal{E}} p_e$. When a node u is selected into the seed set we remove all nodes in its reachability set from the reachability sets of all other nodes. The removal cost can be “charged” to the initial reachability computation.

The dependence of the computation time on the graph size can be improved by using combined reachability sketches [10, 14, 15, 11] instead of maintaining the reachability sets explicitly (see Section 5.1). The sketch size needed in order to provide the required accuracy of $O(\epsilon/s)$ (as in Theorem 10) uniformly for all subsets of size at most s is $k = O(\epsilon^{-2}s^3 \ln n)$. We compute a sketch for each node in each of the r pools, so in total we have rn node sketches. The construction time of these sketches has a term $\sum_{ij} |E_{ij}| = r\ell\bar{m}$ linear in the total size of simulations and a term for sketch constructions which is a product of the number of pools r and the construction time for each pool. The per-pool construction time is as described in Section 5.1 and is bounded by k (sketch size) visits for each node, each involving reverse

traversals of incoming edges of the node in some simulation. The per-pool construction time for an IC model is $O(k(n + \sum_e p_e))$ in expectation. The time with arbitrary simulations for pool i is $O(k(n + \sum_v \max_j d_v(E_{ij})))$. In total over all pools, the construction time is dominated by $O(r(\ell\bar{m} + k(n + m^*)))$, where $m^* = \sum_v \max_{ij} d_v(E_{ij})$ for arbitrary simulations and $m^* = \sum_e p_e$ for simulations generated by an IC model.

The sketches improve the computation time of greedy. Each iteration of greedy uses the (precomputed) union sketch of the current seed set S in each pool. It then examines the sketches of each $v \in V$ to compute the estimate of the averaging oracle $\hat{A}(S \cup \{v\})$ in each pool. This operations takes $O(knr)$ in total for the iteration. Therefore s iterations of greedy maximization takes $O(knrs)$ using the sketches. Combining the construction cost of the sketches using $r = O(s \ln \frac{n}{\delta})$ and $\ell = O(\epsilon^{-2}s^2c)$ and the greedy implementation over the sketches we obtain a total bound on the computation time of

$$\begin{aligned} O(r\ell\bar{m} + kr(m^* + ns)) &= O(r(\ell\bar{m} + k(m^* + ns))) \\ &= O(s \ln \frac{n}{\delta} (\epsilon^{-2}s^2c\bar{m} + \epsilon^{-2}s^3(m^* + ns) \ln n)) \\ &= O(\epsilon^{-2}s^3 \ln \frac{n}{\delta} (c\bar{m} + s(m^* + ns) \ln n)) . \end{aligned} \quad \blacktriangleleft$$

E Optimization with Adaptive Sample Size

The pseudocode for our wrapper is provided in Algorithm 1. The inputs to the wrapper are a base algorithm \mathcal{A} and two constructions of oracles from sets of simulations. The first construction produces an oracle, \hat{F}_v , that we use for validation. The second construction produces oracles, \hat{F}_x , that are provided as input to \mathcal{A} to perform the optimization. The oracles provide an approximation of our influence function $I^\tau(S)$ with non-uniform guarantees. For specified (ϵ, δ) we use the expressions $r_v(\epsilon, \delta)$ or $r_x(\epsilon, \delta)$ for the number of simulations required to obtain (ϵ, δ) guarantees (in the sense of Section 3.5). This gives us a relation between ϵ , δ , and a number of simulations. When constructing an oracle with a given number of simulations r and a specified ϵ , we can determine the confidence δ we have from ϵ and r . The oracles that we consider have the property that for a fixed ϵ , δ decreases at least linearly with the number of simulations. (i.e., when we double the number of simulations δ decreases by at least a factor of 2.)

The wrapper first determines an upper bound ($\lceil \log_2 M/r_x(\epsilon, \delta) \rceil$) on the maximum number of iterations it performs (based on the initial number and the simulation budget) and constructs a validation oracle that provides guarantees for a small number of sets (queries) which equals this maximum number of iterations. It then starts with a set \mathcal{R} of $r_x(\epsilon, \delta)$ simulations that suffice for the oracle \hat{F}_x to provide (non-uniform) (ϵ, δ) approximation guarantees. The wrapper repeats the following: It constructs an “optimization” oracle \hat{F}_x using the set of simulations \mathcal{R} and applies \mathcal{A} over \hat{F}_x to obtain a set T . The wrapper terminates when $\hat{F}_v(T)$ is close to $\hat{F}_x(T)$ or when our simulation budget of M is exceeded. Otherwise, we double the number of simulations in our set \mathcal{R} and repeat.

The wrapped algorithm \mathcal{A} can be an exact or approximate optimizer. It is applied to the oracle function and therefore its quality guarantees are with respect to how well the oracle value $\hat{F}_x(T)$ of the output set T approximates the oracle optimum $\max_{S \parallel |S| \leq s} \hat{F}_x(S)$. The wrapper extends the approximation guarantees that \mathcal{A} provides (with respect to the oracle) to a guarantee with respect to the influence function while avoiding the worst-case number of simulations needed for a uniform approximation.

We first establish some basic properties.

■ **Algorithm 1** Optimization Wrapper.

Input: (i) Two oracle constructions from simulations: \hat{F}_v (validation) and \hat{F}_x (optimization) that with $r_v(\epsilon, \delta)$ (resp., $r_x(\epsilon, \delta)$) simulations provide (ϵ, δ) guarantees. (ii) An optimization algorithm \mathcal{A} that applies to the optimization oracle \hat{F}_x and returns a subset. (iii) M : Bound on maximum number of simulations. (iv) Parameters $\epsilon > 0$ and $\delta > 0$.

```

 $r \leftarrow r_x(\epsilon, \delta)$  // #simulations for  $\hat{F}_x$  to provide  $(\epsilon, \delta)$  guarantees
// Build validation oracle
 $\delta_v \leftarrow \frac{\delta}{\lceil \log_2 M/r \rceil}$ ;  $r_v \leftarrow r_v(\epsilon, \delta_v)$  // #simulations for validation oracle
 $\hat{F}_v \leftarrow$  validation oracle from  $r_v$  i.i.d simulations that provides  $(\epsilon, \delta_v)$  guarantees
 $\mathcal{R} \leftarrow \perp$  // Initialize set of i.i.d simulations for optimization
repeat
  Add  $r$  fresh i.i.d simulations to set  $\mathcal{R}$ 
   $\hat{F}_x \leftarrow$  optimization oracle from simulations  $\mathcal{R}$  that provides  $(\epsilon, *)$  guarantees // * determined
  by  $|\mathcal{R}|$ 
   $T \leftarrow \mathcal{A}(\hat{F}_x)$  // Optimize over the oracle
  if  $\hat{F}_v(T) \geq \frac{(1-2\epsilon)}{1+\epsilon} \hat{F}_x(T)$  then
    | return  $T, \hat{F}_v(T)$ 
  else
    |  $r \leftarrow 2r$ 
until  $|\mathcal{R}| + r_v > M$ 

```

► **Lemma 25.** *Let S be a set with maximum influence (with $I^\tau(S) = \text{OPT}_s^\tau$). With probability at least $1 - 2\delta$, all the optimization oracles \hat{F}_x constructed by the wrapper have $(1 - \epsilon)\text{OPT}_s^\tau \leq \hat{F}_x(S) \leq (1 + \epsilon)\text{OPT}_s^\tau$.*

Proof. The probability that $(1 - \epsilon)\text{OPT}_s^\tau \leq \hat{F}_x(S) \leq (1 + \epsilon)\text{OPT}_s^\tau$ fails for the first oracle is at most δ . The number of \hat{F}_x uses simulations doubles in each iteration and all our constructions are such that the confidence parameter δ decreases at least linearly with the number of simulations. We therefore obtain that the sequence of failure probabilities for $(1 - \epsilon)\text{OPT}_s^\tau \leq \hat{F}_x(S) \leq (1 + \epsilon)\text{OPT}_s^\tau$ is geometric and sums up to at most 2δ . ◀

As an immediate corollary we obtain:

► **Corollary 26.** *Under the conditions of Lemma 25, the oracle optimum in all iterations satisfies*

$$\max_{T \mid |T| \leq s} \hat{F}_x(T) \geq (1 - \epsilon)\text{OPT}_s^\tau .$$

The following is immediate from the construction of the validation oracle.

► **Lemma 27.** *With probability at least $1 - \delta$, the validation oracle has relative error at most ϵ on all tests in which the input set T is such that $I^\tau(T) \geq \text{OPT}_1^\tau$ and absolute error at most $\epsilon\text{OPT}_1^\tau$ otherwise.*

Proof. The wrapper performs at most $\lceil \log_2 M/r \rceil$ iterations before it stops, in each iteration the validation oracle fails to provide an ϵ -approximation with probability at most δ_v . Therefore, by union bound, the probability that the algorithm fails to provide an ϵ -approximation in at least one round is at most $\delta_v \lceil \log_2 M/r \rceil \leq \delta$. ◀

29:30 Sample Complexity Bounds for Influence Maximization

► **Lemma 28.** *Assume that our data and our optimization oracle with r or more simulations, are such that with probability at least $1 - \delta$, the optimum of the oracle is an approximate optimizer, that is:*

$$(1 + \epsilon)\text{OPT}_s^\tau \geq \max_{S \parallel |S| \leq s} \hat{F}_x(S) \geq (1 - \epsilon)\text{OPT}_s^\tau \quad (33)$$

$$I^\tau(\arg \max_{S \parallel |S| \leq s} \hat{F}_x(S)) \geq (1 - \epsilon)\text{OPT}_s^\tau \quad (34)$$

and assume that the algorithm \mathcal{A} returns the oracle optimum. Then with probability at least $1 - 5\delta$, the wrapper terminates after at most $2 \max\{r, r_x(\epsilon, \delta)\} + r_v$ simulations and returns T such that $I^\tau(T) \geq (1 - 5\epsilon)\text{OPT}_s^\tau$.

Proof. First we show that the wrapper returns a set T with the required properties with probability at most $1 - 3\delta$ and then we show that the number of iterations the wrapper does before it stops is smaller than M with probability of at most $1 - 2\delta$.

From Lemma 25, with probability at least $1 - 2\delta$ in all iterations \mathcal{A} returns T for which $\hat{F}_x(T) \geq (1 - \epsilon)\text{OPT}_s^\tau$. The validation succeeds only if $\hat{F}_v(T) \geq \frac{(1-2\epsilon)}{1+\epsilon} \hat{F}_x(T) \geq \frac{(1-\epsilon)(1-2\epsilon)}{1+\epsilon} \text{OPT}_s^\tau$. From Lemma 27 with probability at least $1 - \delta$ in all iterations we have

$$\hat{F}_v(T) \leq \max\{(1 + \epsilon)I^\tau(T), I^\tau(T) + \epsilon\text{OPT}_1^\tau\}.$$

Therefore, with probability $1 - 3\delta$ the set T returned by the wrapper satisfies

$$\frac{(1 - \epsilon)(1 - 2\epsilon)}{1 + \epsilon} \text{OPT}_s^\tau \leq \max\{(1 + \epsilon)I^\tau(T), I^\tau(T) + \epsilon\text{OPT}_1^\tau\}.$$

If $(1 + \epsilon)I^\tau(T) > I^\tau(T) + \epsilon\text{OPT}_1^\tau$ then $I^\tau(T) \geq \frac{(1-\epsilon)(1-2\epsilon)}{(1+\epsilon)^2} \text{OPT}_s^\tau \geq (1 - 5\epsilon)\text{OPT}_s^\tau$. Otherwise, we have that $I^\tau(T) \geq \left(\frac{(1-\epsilon)(1-2\epsilon)}{1+\epsilon} - \epsilon\right) \text{OPT}_s^\tau \geq (1 - 5\epsilon)\text{OPT}_s^\tau$.

We have to show that with probability at least $1 - 2\delta$ within $2 \max\{r, r_x(\epsilon, \delta)\} + r_v$ simulations the wrapper returns such a set T to finish the proof. Consider the first iteration where $|\mathcal{R}| \geq r$. By Equations (33) and (34) with probability at least $1 - \delta$ we have that $I^\tau(T) \geq (1 - \epsilon)\text{OPT}_s^\tau$ and $(1 + \epsilon)\text{OPT}_s^\tau \geq \hat{F}_x(T) \geq (1 - \epsilon)\text{OPT}_s^\tau$. By Lemma 27 we have that with probability at least $1 - \delta$, the validation oracle satisfies that $\hat{F}_v(T) \geq (1 - \epsilon)I^\tau(T)$ or $\hat{F}_v(T) \geq I^\tau(T) - \epsilon\text{OPT}_s^\tau \geq I^\tau(T) - \epsilon\text{OPT}_s^\tau$. By the last two statements we have that with probability of at least $1 - 2\delta$:

$$\hat{F}_v(T) \geq (1 - \epsilon)I^\tau(T) \geq (1 - \epsilon)^2 \text{OPT}_s^\tau \geq (1 - 2\epsilon)\text{OPT}_s^\tau \geq \frac{(1 - 2\epsilon)}{1 + \epsilon} \hat{F}_x(T)$$

or

$$\hat{F}_v(T) \geq I^\tau(T) - \epsilon\text{OPT}_s^\tau \geq (1 - 2\epsilon)\text{OPT}_s^\tau \geq \frac{(1 - 2\epsilon)}{1 + \epsilon} \hat{F}_x(T). \quad \blacktriangleleft$$

Theorem 8, which we restate below to provide reading fluency, now follows as a corollary.

► **Theorem 29** (Theorem 8). *Suppose that on our data the averaging (respectively, median-of-averages) oracle \hat{F} has the property that with r simulations, with probability at least $1 - \delta$, the oracle optimum $T := \arg \max_{S \parallel |S| \leq s} \hat{F}(S)$ satisfies*

$$I^\tau(T) \geq (1 - \epsilon)\text{OPT}_s^\tau.$$

Then with probability at least $1 - 5\delta$, when using $2 \max\{r, r(\epsilon, \delta)\} + O(\epsilon^{-2}c(\ln \frac{1}{\delta} + \ln(\ln \ln \frac{n}{\delta} + \ln s)))$ simulations with the median-of-averages oracle and $2 \max\{r, r(\epsilon, \delta)\} + O(\epsilon^{-2}c(\ln \frac{1}{\delta} + \ln(\ln \ln \frac{n}{\delta} + \ln n)))$ simulations with the averaging oracle, the wrapper outputs a set T such that $I^\tau(T) \geq (1 - 5\epsilon)\text{OPT}_s^\tau$.

Proof of Theorem 8. We analyze here the number of simulations required using the averaging oracles and the median-of-averages oracles, in both cases we use median-of-averages oracles for validation. In both cases $r_v = r(\epsilon, \delta_v) = O(\epsilon^{-2}c \log \frac{1}{\delta_v})$, where $\delta_v = \frac{\delta}{\lceil \log_2 \frac{M}{r_x} \rceil}$. By Lemma 28 the number of simulations is at most $2 \max\{r, r_x(\epsilon, \delta)\} + r_v$. M and r_x get different values for each oracle.

Median-of-averages oracles analysis. We have that $r_x = O(\epsilon^{-2}c \ln \delta^{-1})$ by Lemma 6 and we set $M = O(\epsilon^{-2}cs \ln \frac{n}{\delta})$ by Theorem 7. Simple calculation shows that:

$$\frac{1}{\delta_v} = \frac{\lceil \ln \frac{s \ln \frac{n}{\delta}}{\ln \frac{1}{\delta}} \rceil}{\delta} \leq \frac{1}{\delta} \left(\ln s + \ln \left(\ln \frac{n}{\delta} \right) \right).$$

Therefore,

$$r_v = O \left(\epsilon^{-2}c \left(\ln \frac{1}{\delta} + \ln \left(\ln \ln \frac{n}{\delta} + \ln s \right) \right) \right).$$

Averaging oracles analysis. We have $r_x = O(\epsilon^{-2}c\delta^{-1})$ and we set $M = O(\epsilon^{-2}sn \ln \frac{n}{\delta})$ according to the respective worst-case guarantees on the number of simulations specified in (1). A simple calculations shows:

$$\frac{1}{\delta_v} = \frac{\lceil \ln \frac{\delta sn \ln \frac{n}{\delta}}{c} \rceil}{\delta} \leq \frac{1}{\delta} \left(\ln \ln \frac{n}{\delta} + 2 \ln n \right)$$

Therefore,

$$r_v = O \left(\epsilon^{-2}c \left(\ln \frac{1}{\delta} + \ln \left(\ln \ln \frac{n}{\delta} + \ln n \right) \right) \right). \quad \blacktriangleleft$$

We next consider cases where the algorithm \mathcal{A} is approximate (may not return the oracle optimizer). We assume in these cases that the optimization oracles \hat{F}_x when constructed with a given number of simulations provide, with high probability, uniform ϵ -approximation for all $\binom{n}{s}$ subsets of cardinality at most s :

$$\forall T \text{ such that } |T| < s, \quad \left| \hat{F}(T) - I^\tau(T) \right| \leq \epsilon \max\{I^\tau(T), \text{OPT}_1^\tau\}.$$

We first show that a very weak assumption on \mathcal{A} suffices to guarantee termination with good probability.

► **Lemma 30.** *If the optimization oracle \hat{F}_x when constructed with r or more simulations provides uniform ϵ -approximation with probability at least $1 - \delta$, and the algorithm \mathcal{A} returns T such that $\hat{F}_x(T) \geq (1 - \epsilon)\text{OPT}_1^\tau$. Then with probability at least $1 - 2\delta$ the wrapper will terminate after using at most $2 \max\{r, r_x(\epsilon, \delta)\} + r_v$ simulations.*

Proof. Consider the first iteration where \hat{F}_x is constructed using at least r simulations. Let T be the set that \mathcal{A} returns at this iteration. Since \hat{F}_x provides uniform ϵ -approximation we have that $\hat{F}_x(T) \leq (1 + \epsilon)I^\tau(T)$ with probability at least $1 - \delta$. Combining this with our assumption we get that $I^\tau(T) \geq \text{OPT}_1^\tau$ and by Lemma 27 we have that with probability at least $1 - \delta$ if $I^\tau(T) \geq \text{OPT}_1^\tau$ then $\hat{F}(T) \geq (1 - \epsilon)I^\tau(T)$ and if $\frac{1-\epsilon}{1+\epsilon}\text{OPT}_1^\tau \leq I^\tau(T) \leq \text{OPT}_1^\tau$ then $\hat{F}_v(T) \geq I^\tau(T) - \epsilon\text{OPT}_1^\tau \geq I^\tau(T) - \frac{\epsilon(1+\epsilon)}{1-\epsilon}I^\tau(T)$. Combining we obtain that $\hat{F}_v(T) \geq \frac{1-2\epsilon}{1+\epsilon}\hat{F}_x(T)$, and thus the validation condition holds. ◀

29:32 Sample Complexity Bounds for Influence Maximization

We next consider algorithms \mathcal{A} that guarantees some approximation ratio ρ .

► **Theorem 31.** *Suppose that our optimization oracle when constructed with r or more simulations provides uniform ϵ -approximation with probability at least $1 - \delta$. Assume now that the algorithm \mathcal{A} returns a set T such that*

$$\hat{F}_x(T) \geq \rho \max_{S \parallel |S| \leq s} \hat{F}_x(S) .$$

Then, the set T returned by our wrapper satisfies $\mathcal{I}^\tau(T) \geq \rho(1 - 5\epsilon)\text{OPT}_s^\tau$ with probability of at least $(1 - 3\delta)$.

Proof. Consider an optimal set S (with $\mathcal{I}^\tau(S) = \text{OPT}_s^\tau$). By Lemma 25 with probability at least $1 - 2\delta$ all our oracles have $(1 - \epsilon)\text{OPT}_s^\tau \leq \hat{F}_x(S) \leq (1 + \epsilon)\text{OPT}_s^\tau$ are within $(1 \pm \epsilon)\text{OPT}_s^\tau$. By the assumption, the sets T returned by \mathcal{A} in all iterations have $\hat{F}_x(T) \geq \rho \max_{S \parallel |S| \leq s} \hat{F}_x(S) \geq \rho(1 - \epsilon)\text{OPT}_s^\tau$. When the wrapper stops we have that $\hat{F}_x(T) \leq \frac{1+\epsilon}{1-2\epsilon} \hat{F}_v(T)$ and by Lemma 27 we have with probability at least $1 - \delta$ that $\hat{F}_v(T) \leq \max\{(1 + \epsilon)\mathcal{I}^\tau(T), \mathcal{I}^\tau(T) + \epsilon\text{OPT}_1^\tau\}$.

Combining, we have that with probability at least $1 - 3\delta$,

$$\rho(1 - \epsilon)\text{OPT}_s^\tau \leq \rho \hat{F}_x(S) \leq \hat{F}_x(T) \leq \frac{1 + \epsilon}{1 - 2\epsilon} \hat{F}_v(T) \leq \frac{1 + \epsilon}{1 - 2\epsilon} \max\{(1 - \epsilon)\mathcal{I}^\tau(T), \mathcal{I}^\tau(T) + \epsilon\text{OPT}_1^\tau\}.$$

Now, a simple calculation shows that $\mathcal{I}^\tau(T) \geq \rho(1 - 5\epsilon)\text{OPT}_s^\tau$. ◀

We can prove now Theorem 9 (restated for reading fluency):

► **Theorem 32** (Theorem 9). *If the averaging oracle \hat{A} has the property that with $\geq r$ simulations, with probability at least $1 - \delta$, it provides a uniform ϵ -approximation for all subsets of size at most s , then with $2 \max\{r, r(\epsilon, \delta)\} + O(\epsilon^{-2}c(\ln \frac{1}{\delta} + \ln(\ln \ln \frac{n}{\delta} + \ln n)))$ simulations we can find in polynomial time a $(1 - (1 - 1/s)^s)(1 - 5\epsilon)$ approximate solution with confidence $1 - 5\delta$.*

Proof. The averaging oracle is monotone and submodular [31] and therefore greedy can efficiently recover a set T such that $\hat{F}_x(T) \geq (1 - (1 - 1/s)^s) \max_{S \parallel |S| \leq s} \hat{F}_x(S)$.

By Lemma 30, the wrapper terminates using at most $2 \max\{r, r_x(\epsilon, \delta)\} + r_v$ with probability at least $1 - 2\delta$. Applying Theorem 31 with $\rho = (1 - (1 - 1/s)^s)$, we get that $\mathcal{I}^\tau(T) \geq (1 - (1 - 1/s)^s)(1 - 5\epsilon)\text{OPT}_s^\tau$ with probability at least $1 - 3\delta$. Hence, with probability at least $1 - 5\delta$ the wrapper applied with greedy finds $(1 - (1 - 1/s)^s)(1 - 5\epsilon)$ -approximate solution using $2 \max\{r, r_x(\epsilon, \delta)\} + r_v$ simulations. ◀

F Variance Bounds for Dependent Models

In this section we provide a proof for Corollary 3. We consider a natural extensions of IC models, b -dependence, that allow for some dependencies between edges and mixtures of IC and IGT models. For these extensions, we establish upper bounds of the form (8) on the variance of the reachability of a set of nodes.

We bound the variance by constructing for each dependent model a corresponding IC model and then apply the variance upper bound established in Section A for IC models.

For mixture models we provide a generic derivation that bounds the variance of the mixture by variance of components.

F.1 b -Dependence Models

The first family we consider are b -dependence models, which we define as follows. We assume that all edges with the same tail node are partitioned into disjoint groups where each group is of size at most b . The edges of each group B are either all active together with probability p_B or none is active with probability $1 - p_B$. The special case where all groups are of size 1 corresponds to an IC model (where all edges are independent).

► **Theorem 33.** *Let \mathcal{G} be a b -dependence model for some $b \geq 1$. For every set T we have that:*

$$\text{Var}[R_{\mathcal{G}}^{\tau}(T)] \leq 2b\tau I_{\mathcal{G}}^{\tau}(T) \max_{v \in V \setminus T} I_{\mathcal{G}}^{\tau}(v)$$

Proof. We construct an IC model \mathcal{G}' from the given b -dependence model \mathcal{G} . The model \mathcal{G}' is defined over the set of nodes V of \mathcal{G} together with an additional set D of dummy nodes. The construction has the properties that 2τ -step influence in \mathcal{G}' from a set of nodes $T \subseteq V$ is equal to τ -stepped influence of T in \mathcal{G} . Furthermore, the variances of the sizes of the 2τ -step reachability of T in \mathcal{G}' is the same as the variance of the τ step reachability of T in \mathcal{G} . The influence of each dummy node in \mathcal{G}' is at most $b \max_{v \in V} I^{\tau}(v)$. The claim follows from these properties and Theorem 2.

Here is a formal description of our reduction.

We start by putting in \mathcal{G}' the set V of the nodes of \mathcal{G} . Then for every group $B = \{(u, v_1), (u, v_2), \dots, (u, v_{\ell})\}$ in \mathcal{G} we do the following:

1. Add a new dummy node v_B to \mathcal{G}' , and add to \mathcal{G}' the edge (u, v_B) and give it the probability p_B . We assign weight 0 to v_B so that it does not contribute to the reachability of any set of nodes.
2. we create edges (v_B, v_i) for every $1 \leq i \leq \ell$, each such edge has probability 1.

Let $T \subset V$ be a set of nodes in \mathcal{G} . It follows from our construction that for any set of nodes $B \subset V$ the probability that $R_{\mathcal{G}'}^{2\tau}(T) = B$ is the same as the probability that $R_{\mathcal{G}}^{\tau}(T) = B$. This implies that for any $T \subseteq V$

$$I_{\mathcal{G}'}^{2\tau}(T) = I_{\mathcal{G}}^{\tau}(T) ,$$

and

$$\text{Var}[R_{\mathcal{G}'}^{2\tau}(T)] = \text{Var}[R_{\mathcal{G}}^{\tau}(T)] .$$

Each dummy node is connected to at most k original nodes, hence, $I_{\mathcal{G}'}^{2\tau}(v)$ is bounded by $b \max_{v \in V} I_{\mathcal{G}}^{2\tau}(v)$. By Theorem 20 it follows that for every set of nodes T in \mathcal{G}' :

$$\text{Var}[R_{\mathcal{G}'}^{2\tau}(T)] \leq 2\tau I_{\mathcal{G}'}^{2\tau}(T) \max_{v \in V \setminus T} I_{\mathcal{G}'}^{2\tau}(v).$$

Combining all these observations together, we get that

$$\text{Var}[R_{\mathcal{G}}^{\tau}(T)] \leq 2b\tau I_{\mathcal{G}}^{\tau}(T) \max_{v \in V \setminus T} I_{\mathcal{G}}^{\tau}(v) \quad \blacktriangleleft$$

This Theorem can be generalized to more complex dependencies. For example it holds for any distribution on subsets of the outgoing edges from each node that we can realize by a distribution on disjoint subsets where we draw each subset with certain probability, and take the union of the subset which we draw.

F.2 Mixture of IC and IGT Models

The second family of dependent models we consider is a mixture of IC and IGT models.

Consider a set of models $\mathcal{G}_i(V)$ for $i \in [r]$ and respective probabilities p_i such that $\sum_{i=1}^r p_i = 1$. We define a mixture model $\mathcal{G}(V)$ as follows. To draw $\phi \sim \mathcal{G}$, we first draw $i \in [r]$ according to probabilities p_i and then return $\phi \sim \mathcal{G}_i$.

We provide two proofs for the variance bound of the mixture. The first is direct and applies to any mixture of models that satisfies the variance bound of Theorem 2), and in particular to mixtures of strongly submodular SDMs. The second proof is specific to live-edge models and based on a reduction to an IC model.

► **Theorem 34.** *Consider a model \mathcal{G} that is a mixture of r models \mathcal{G}_i with probabilities p_i that satisfy the variance bound of Theorem 2. Then for all $T \subset V$,*

$$\text{Var}[\mathbf{R}_{\mathcal{G}}^{\tau}(T)] \leq \frac{\tau + 1}{\min_i p_i} I_{\mathcal{G}}^{\tau}(T) \max_{v \in V} I_{\mathcal{G}}^{\tau}(v)$$

Proof. We first relate the influence of T in the mixture model to the influence of T in the components.

$$\mathbf{I}_{\mathcal{G}}^{\tau}(T) = \mathbb{E}[\mathbf{R}_{\mathcal{G}}^{\tau}(T)] = \sum_{i=1}^r p_i \mathbb{E}[\mathbf{R}_{\mathcal{G}_i}^{\tau}(T)] = \sum_{i=1}^r p_i \mathbf{I}_{\mathcal{G}_i}^{\tau}(T). \quad (35)$$

This holds to any set T and any τ . Therefore we also obtain the inequality

$$\mathbf{M}_{\mathcal{G}}^{\tau}(\bar{T}) = \max_{v \in V \setminus T} I_{\mathcal{G}}^{\tau}(v) = \max_{v \in V \setminus T} \sum_{i=1}^r p_i I_{\mathcal{G}_i}^{\tau}(v) \leq \sum_{i=1}^r p_i \max_{v \in V \setminus T} I_{\mathcal{G}_i}^{\tau}(v) = \sum_{i=1}^r p_i \mathbf{M}_{\mathcal{G}_i}^{\tau}(\bar{T}). \quad (36)$$

It also follows that we can bound the influence values on the component by the respective ones in the mixture: $\mathbf{I}_{\mathcal{G}_i}^{\tau}(T) \leq \frac{1}{p_i} \mathbf{I}_{\mathcal{G}}^{\tau}(T)$ and thus

$$\max_i \mathbf{I}_{\mathcal{G}_i}^{\tau}(T) \leq \frac{1}{\min_i p_i} \mathbf{I}_{\mathcal{G}}^{\tau}(T) \quad (37)$$

$$\max_i \mathbf{M}_{\mathcal{G}_i}^{\tau}(\bar{T}) \leq \frac{1}{\min_i p_i} \mathbf{M}_{\mathcal{G}}^{\tau}(\bar{T}). \quad (38)$$

The random variable $\mathbf{R}_{\mathcal{G}}^{\tau}(T)$ can be expressed as a sum of r products of random variables:

$$\mathbf{R}_{\mathcal{G}}^{\tau}(T) = \sum_{i=1}^r X_i \mathbf{R}_{\mathcal{G}_i}^{\tau}(T),$$

where X_i are Bernoulli with probabilities p_i . The random variables $\{\mathbf{R}_{\mathcal{G}_i}^{\tau}(T)\}$ are independent of each other and also are independent from (the joint distribution of) $\{X_i\}$. The variables $\{X_i\}$ have negative dependence as $\sum_i X_i = 1$ and thus the products $X_i \mathbf{R}_{\mathcal{G}_i}^{\tau}(T)$ are also negatively dependent and thus

$$\text{Var}[\mathbf{R}_{\mathcal{G}}^{\tau}(T)] \leq \sum_i \text{Var}[X_i \mathbf{R}_{\mathcal{G}_i}^{\tau}(T)].$$

We will instead bound the variance of a surrogate random variable

$$Y = \sum_i X_i \mathbf{R}_{\mathcal{G}_i}^{\tau}(T)$$

that has the same sum of products but with the variables $\{X_i\}$ being independent of each other and hence the products are also independent. We have

$$\text{Var}[Y] = \sum_i \text{Var}[X_i \mathbf{R}_{\mathcal{G}_i}^\tau(T)] \geq \text{Var}[\mathbf{R}_{\mathcal{G}}^\tau(T)] \quad (39)$$

We next express the variance of each product using variance properties of the product of two independent random variables. For $i \in [r]$:

$$\begin{aligned} \text{Var}[X_i \mathbf{R}_{\mathcal{G}_i}^\tau(T)] &= \text{Var}[X_i] \text{Var}[\mathbf{R}_{\mathcal{G}_i}^\tau(T)] + \mathbb{E}[X_i]^2 \text{Var}[\mathbf{R}_{\mathcal{G}_i}^\tau(T)]^2 + \text{Var}[X_i] \mathbb{E}[\mathbf{R}_{\mathcal{G}_i}^\tau(T)]^2 \\ &= p_i(1-p_i) \text{Var}[\mathbf{R}_{\mathcal{G}_i}^\tau(T)] + p_i^2 \text{Var}[\mathbf{R}_{\mathcal{G}_i}^\tau(T)] + p_i(1-p_i) \mathbf{I}_{\mathcal{G}_i}^\tau(T)^2 \\ &= p_i \text{Var}[\mathbf{R}_{\mathcal{G}_i}^\tau(T)] + p_i(1-p_i) \mathbf{I}_{\mathcal{G}_i}^\tau(T)^2. \end{aligned}$$

Therefore, invoking Theorem 2 to bound the variance for each IC model \mathcal{G}_i and then using (37) and (38) and finally using (35) and (36) we get

$$\begin{aligned} \text{Var}[Y] &= \sum_{i=1}^r p_i (\text{Var}[\mathbf{R}_{\mathcal{G}_i}^\tau(T)] + p_i(1-p_i) \mathbf{I}_{\mathcal{G}_i}^\tau(T)^2) \\ &\leq \sum_{i=1}^r p_i \text{Var}[\mathbf{R}_{\mathcal{G}_i}^\tau(T)] + \sum_{i=1}^r p_i \mathbf{I}_{\mathcal{G}_i}^\tau(T)^2 \\ &\stackrel{\text{Theorem 2}}{\leq} \sum_{i=1}^r p_i \tau \mathbf{M}_{\mathcal{G}_i}^{\tau-1}(\bar{T}) \mathbf{I}_{\mathcal{G}_i}^\tau(T) + \sum_{i=1}^r p_i \mathbf{I}_{\mathcal{G}_i}^\tau(T)^2 \\ &\stackrel{(37), (38)}{\leq} \frac{\tau}{\min_i p_i} \mathbf{I}_{\mathcal{G}}^\tau(T) \sum_{i=1}^r p_i \mathbf{M}_{\mathcal{G}_i}^{\tau-1}(\bar{T}) + \frac{1}{\min_i p_i} \mathbf{I}_{\mathcal{G}}^\tau(T) \sum_{i=1}^r p_i \mathbf{I}_{\mathcal{G}_i}^\tau(T) \\ &\stackrel{(35), (36)}{\leq} \frac{1}{\min_i p_i} \mathbf{I}_{\mathcal{G}}^\tau(T) (\tau \mathbf{M}_{\mathcal{G}}^{\tau-1}(\bar{T}) + \mathbf{I}_{\mathcal{G}}^\tau(T)) \leq \frac{\tau+1}{\min_i p_i} \mathbf{I}_{\mathcal{G}}^\tau(T) \max_{v \in V} \mathbf{I}_{\mathcal{G}}^\tau(v) \end{aligned}$$

◀

We next give a different proof (of a slightly different bound) for live-edge models using a reduction to an IC model. Consider a set of τ -steps models $\mathcal{G}_i(V, \mathcal{E}_i)$ for $i \in [r]$ and respective probabilities p_i such that $\sum_{i=1}^r p_i = 1$. We define a mixture model $\mathcal{G}(V, \bigcup_i \mathcal{E}_i)$ as follows. To draw $E \sim \mathcal{G}$, we first draw $i \in [r]$ according to probabilities p_i and then return $E \sim \mathcal{G}_i$.

► **Theorem 35.** *Consider a model \mathcal{G} that is a mixture of r IC models \mathcal{G}_i with probabilities p_i . Then for all $T \subset V$,*

$$\text{Var}[\mathbf{R}_{\mathcal{G}}^\tau(T)] \leq \frac{\tau+1}{\min_i p_i} \mathbf{I}_{\mathcal{G}}^\tau(T) \max\{\mathbf{I}_{\mathcal{G}}^\tau(T), \max_{v \in V} \mathbf{I}_{\mathcal{G}}^\tau(v)\}$$

Proof. We first argue that we can assume without loss of generality that T is a single node and $\bigcup_i \mathcal{E}_i$ does not contain edges that are incoming to T . We can transform a general case \mathcal{G} and T to this form by contracting all nodes in T into a single node and deleting all edges that are incoming to T . We then retain the same conditional distribution on the remaining edges. Note that this transformation preserves the distribution of $\mathbf{R}_{\mathcal{G}}^\tau(T)$ and hence also its expectation and variance. The influence values $\mathbf{I}_{\mathcal{G}}^\tau(v)$ of nodes $v \in V \setminus T$ can only decrease. Finally, the transformed model is also a mixture of correspondingly transformed IC models, where in each such model the distribution of $\mathbf{R}_{\mathcal{G}_i}^\tau(T)$ remains the same and influence values $\mathbf{I}_{\mathcal{G}_i}^\tau(v)$ can only decrease. It follows that the claimed variance bound for the transformed model implies the same bound for the original model.

29:36 Sample Complexity Bounds for Influence Maximization

We construct a new IC model \mathcal{G}' with respect to (a single node) T as follows. The new model has nodes $V' = \{v\} \cup \bigcup_i V_i$, where each V_i is a map of V . We create an instantiation of each of our IC models \mathcal{G}_i with set of nodes V_i and edges \mathcal{E}_i with the probabilities as in the model \mathcal{G}_i . The new IC model \mathcal{G}' has a root node v with weight 0 and for each $i \in [r]$, there is an edge (v, T_i) with probability p_i , where T_i is the image of T in the copy of \mathcal{G}_i . We can see that

$$I_{\mathcal{G}'}^{\tau+1}(v) = I_{\mathcal{G}}^{\tau}(T) = \sum_{i=1}^r p_i I_{\mathcal{G}_i}^{\tau}(T) , \quad (40)$$

that is, the $\tau + 1$ steps influence of v in the constructed IC model \mathcal{G}' is equal to the τ steps influence of T in the mixture model \mathcal{G} .

We next consider the variance of the random variables $\mathbf{R}_{\mathcal{G}}^{\tau}(T)$ and $\mathbf{R}_{\mathcal{G}'}^{\tau+1}(v)$. Both these random variables are a sum of r products of random variables:

$$\sum_{i=1}^r X_i \mathbf{R}_{\mathcal{G}_i}^{\tau}(T) ,$$

where X_i are Bernoulli with probabilities p_i . In both cases the random variables $\{\mathbf{R}_{\mathcal{G}_i}^{\tau}(T)\}$ are independent of each other and also are independent from (the joint distribution of) $\{X_i\}$. But in the case of $\mathbf{R}_{\mathcal{G}'}^{\tau+1}(v)$ the random variables X_i are independent and hence also the products are independent and in the case of $\mathbf{R}_{\mathcal{G}}^{\tau}(T)$, the variables $\{X_i\}$ have negative dependence as $\sum_i X_i = 1$ and thus the products $X_i \mathbf{R}_{\mathcal{G}_i}^{\tau}(T)$ are also negatively dependent. Therefore,

$$\text{Var}[\mathbf{R}_{\mathcal{G}}^{\tau}(T)] \leq \text{Var}[\mathbf{R}_{\mathcal{G}'}^{\tau+1}(v)] = \sum_i \text{Var}[X_i \mathbf{R}_{\mathcal{G}_i}^{\tau}(T)] . \quad (41)$$

Finally, we bound $\mathbf{M}_{\mathcal{G}'}^{\tau}(\bar{v})$ by considering the maximum influence of a node other than v in the constructed model \mathcal{G}' . For T_i we have

$$I_{\mathcal{G}'}^{\tau}(T_i) = I_{\mathcal{G}_i}^{\tau}(T) \leq \frac{1}{p_i} I_{\mathcal{G}}^{\tau}(T) , \quad (42)$$

where the last inequality follows from (40). We next consider nodes $z_i \in V_i$ that is a map of a node $z \in V$.

$$I_{\mathcal{G}'}^{\tau}(z_i) = I_{\mathcal{G}_i}^{\tau}(z) \leq \frac{1}{p_i} I_{\mathcal{G}}^{\tau}(z) . \quad (43)$$

The last inequality follows because for any node $z \in v$ we have $I_{\mathcal{G}}^{\tau}(z) = \sum_{i=1}^r p_i I_{\mathcal{G}_i}^{\tau}(z)$. Combining (42) and (43) we get

$$\begin{aligned} \mathbf{M}_{\mathcal{G}'}^{\tau}(\bar{v}) &= \max_{u \in V' \setminus \{v\}} I_{\mathcal{G}'}^{\tau}(u) \leq \max_{u \in V} \max_i I_{\mathcal{G}_i}^{\tau}(u) \leq \max_{u \in V} \frac{1}{\min_i p_i} I_{\mathcal{G}}^{\tau}(u) \\ &= \frac{1}{\min_i p_i} \max\{I_{\mathcal{G}}^{\tau}(T), \max_{z \in V} I_{\mathcal{G}}^{\tau}(z)\} . \end{aligned} \quad (44)$$

To conclude, we invoke Theorem 2 for the IC model \mathcal{G}' :

$$\text{Var}[\mathbf{R}_{\mathcal{G}}^{\tau}(T)] \stackrel{(41)}{\leq} \text{Var}[\mathbf{R}_{\mathcal{G}'}^{\tau+1}(v)] \stackrel{\text{Theorem 2}}{\leq} (\tau + 1) I_{\mathcal{G}'}^{\tau+1}(v) \mathbf{M}_{\mathcal{G}'}^{\tau}(\bar{v}) .$$

We then apply inequalities (44) and the equality (40) to obtain the claim. \blacktriangleleft