Adaptive Influence Maximization with Myopic Feedback

Binghui Peng[1], Wei Chen[2]
[1]Columbia University
[2]Microsoft Research

Influence Maximization

Input: a social network $G(V, E, p)$ in a stochastic diffusion model, a budget k.
Output: k seed nodes with the largest expected *influence spread*.
Applications: viral marketing, rumor control, etc.

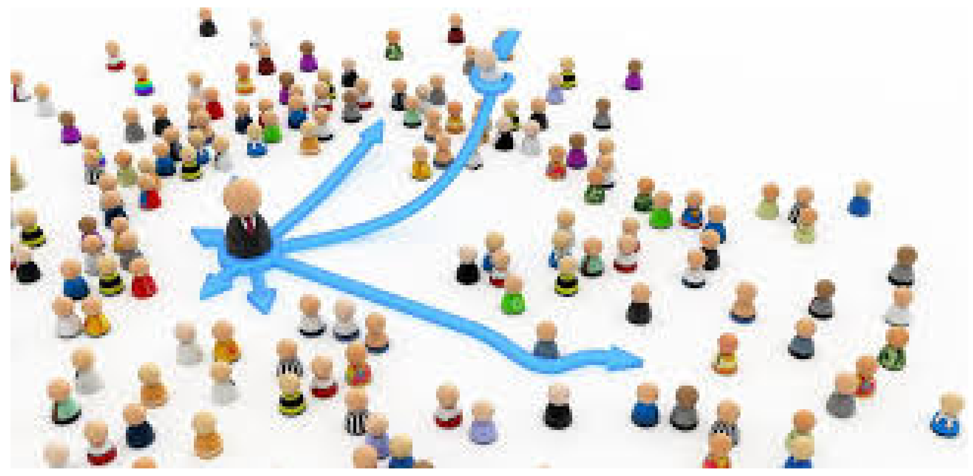

Diffusion Model

Independent Cascade (IC) model

Each node u has an independent chance p_{uv} to pass information to its neighbour node v.

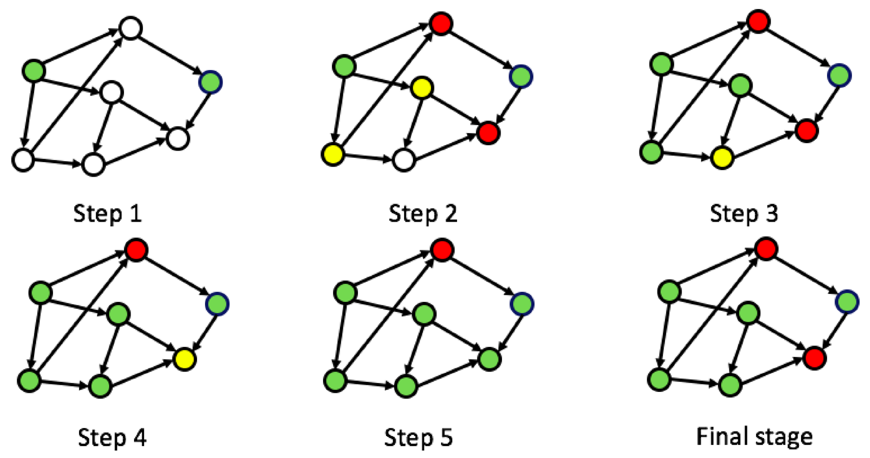

Figure 1: Independent Cascade model

Adaptive Influence Maximization

Adaptive Influence Maximization: an *adaptive* algorithm selects seeds one after one, and each selected seed returns feedback information containing the local status.

Feedback model

- **Full-adoption feedback**: feedback information contains the full cascade from the selected seed.
- **Myopic feedback**: feedback information only contains the immediate neighbors of the selected seed.

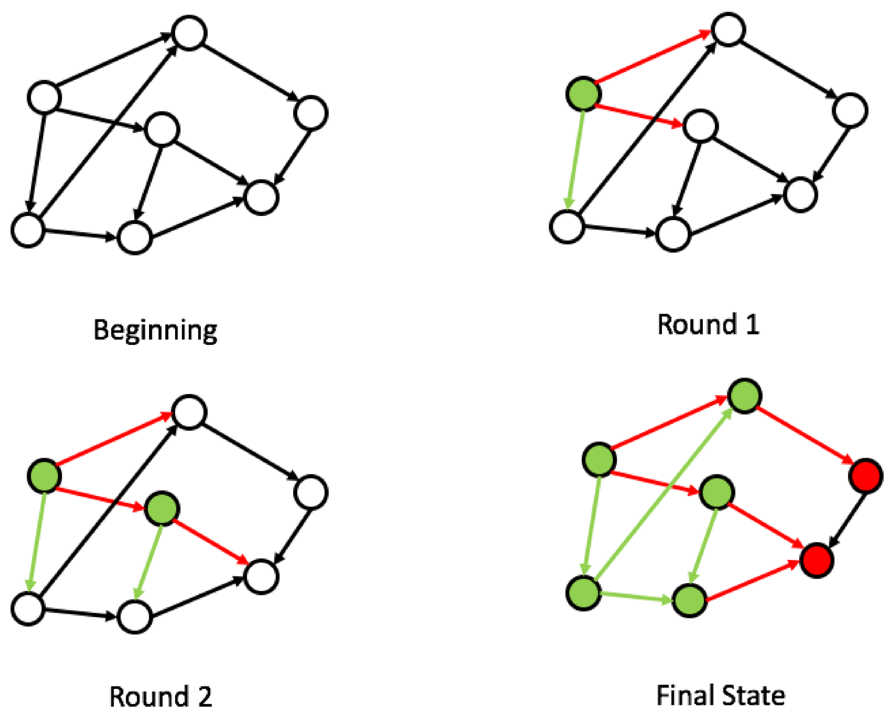

Figure 2: Myopic Feedback

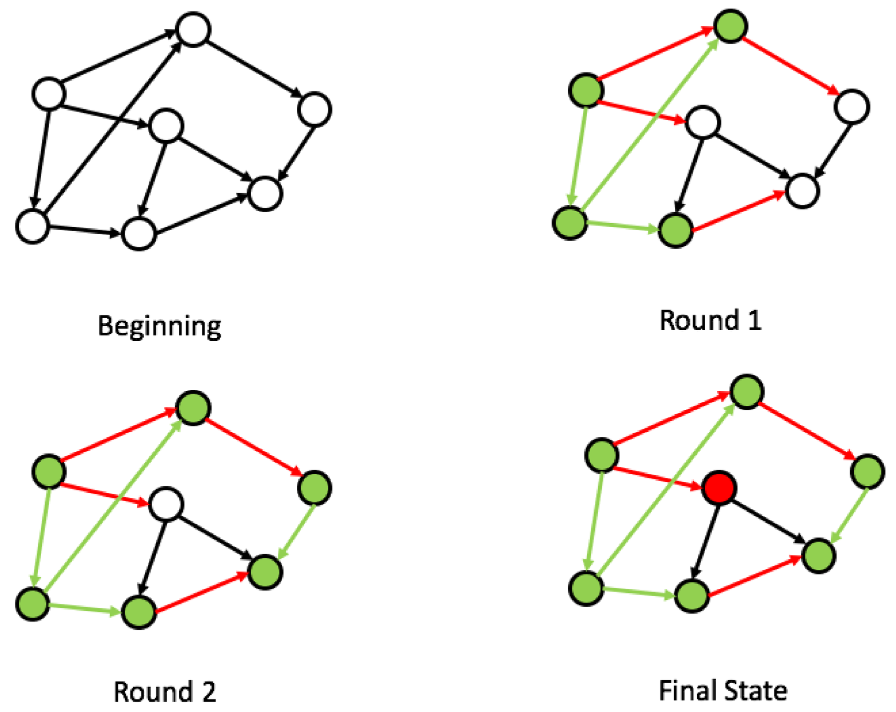

Figure 3: Full-adoption Feedback

Property

Submodularirty

$$\Delta(u \mid A) \geq \Delta(u \mid B), A \subseteq B \subseteq V.$$

Adaptive Submodularity

$$\Delta(u \mid \psi) \geq \Delta(u \mid \psi'), \psi \subseteq \psi'.$$

The influence spread function is submodular under IC model [1], and it is adaptive submodular with full-adoption feedback [2].

Challenge

The influence spread function is **not** adaptive submodular with **myopic feedback**.

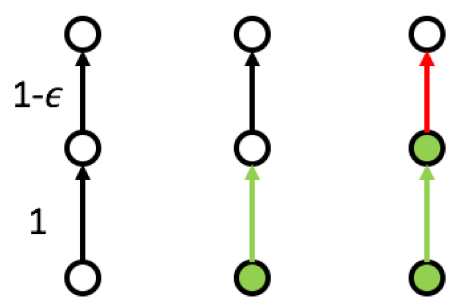

Figure 4: A counter-example

Contribution

Our main contributions are:

- We consider the *adaptivity gap*, i.e., the supremum ratio between the optimal adaptive influence spread and the optimal non-adaptive influence spread, and show that the adaptivity gap is between $[\frac{e}{e-1}, 4]$.
- We show that the approximation ratio of both non-adaptive greedy and adaptive greedy algorithms are in $[\frac{1}{4}(1-\frac{1}{e}), \frac{e^2+1}{(e+1)^2}]$, which confirms an open conjecture of Golvin&Krause(2011) [2].

Overview of Techniques

Upper bound on adaptivity gap:

Idea 1. Compare an adaptive policy π with the random walk non-adaptive policy $\mathcal{W}(\pi)$, which picks a random leaf of the decision tree of π.

Idea 2. Define t-th aggregate influence spread function $\sigma^t(S)$, in which seeds have t chance to activate neighbors.

$$\sigma^t(S) = \mathop{\mathbb{E}}_{\Phi^1,\cdots,\Phi^t \sim \mathcal{P}} \left[f^t(S, \Phi^1, \cdots, \Phi^t) \right].$$

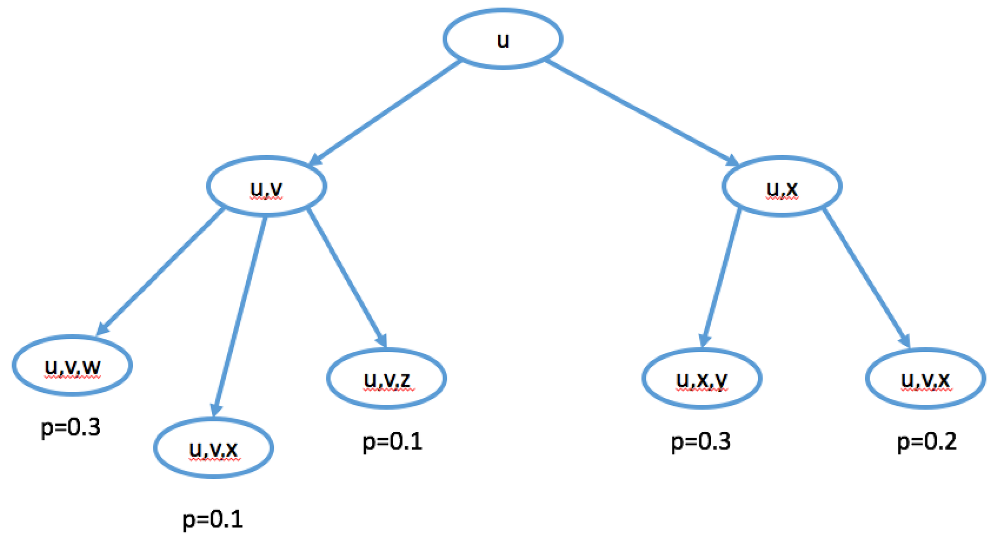

Figure 5: Random walk non-adaptive policy

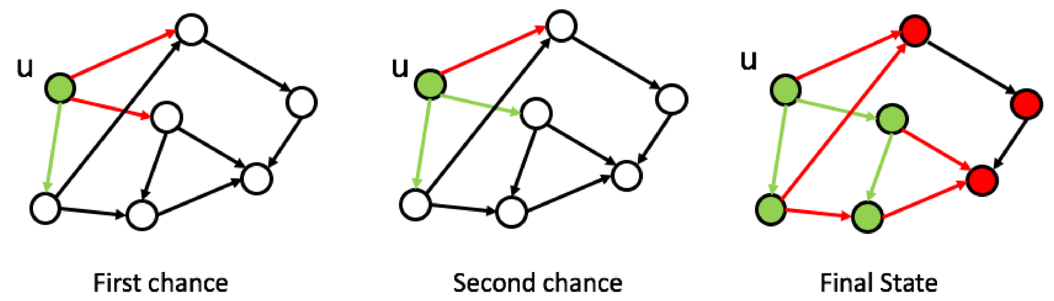

Figure 6: Aggregate influence spread function $\sigma^2(u)$

Idea 3. Construct a fictitious hybrid policy $\bar{\pi}$ which runs in *three* independent realizations of the graphs, use coupling arguments to connect the fictitious hybrid policy $\bar{\pi}$ with the non-adaptive policy π.

Finally. We pin down a constant upper bound via a chain of inequalities,

$$\sigma(\pi) \leq \sigma(\bar{\pi}) = \sigma^3(\pi) \leq 2\sigma^2(\mathcal{W}(\pi)) \leq 4\sigma(\mathcal{W}(\pi)).$$

Future Direction

1. The Adaptivity gap in the full-adoption feedback model is still open.
2. The approximation ratio of (adaptvie) greedy algorithm in the Linear Threstold model is still open.

Reference

[1] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD*, pages 137–146. ACM, 2003.

[2] Daniel Golovin and Andreas Krause. Adaptive submodularity:theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.