Microsoft Research

Microsoft

Influence Maximization: Pushing the Limits of Combinatorial Optimizations and Online Learning

Wei Chen 陈卫
Microsoft Research Asia

Information and Influence Propagation in Networks

Examples of Studies on Influence in Networks: Obesity and Stopping Smoking

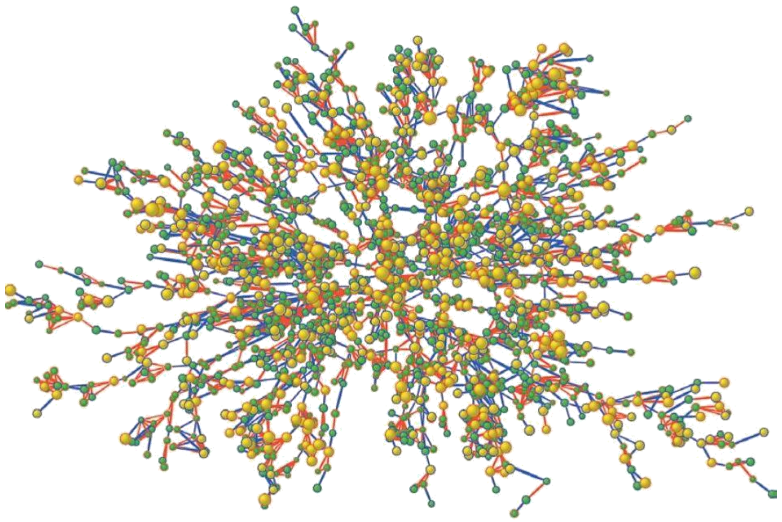

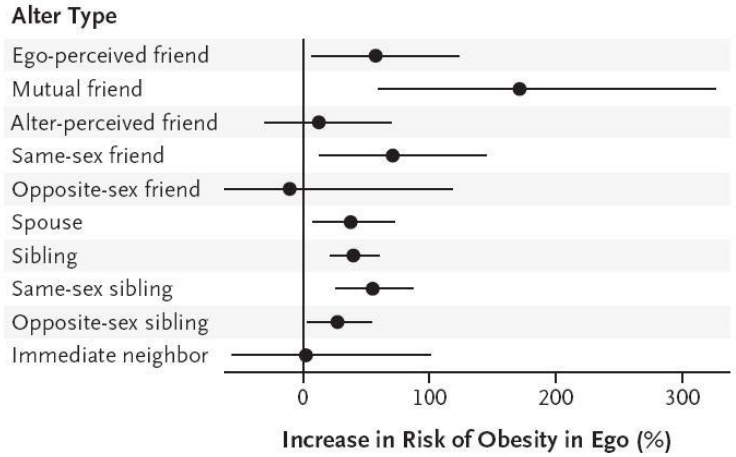

Christakis N A and Fowler J H. The spread of obesity in a large social network over 32 years. New England Journal of Medicine, 2007(357.4):370~379

Christakis N A and Fowler J H. The collective dynamics of smoking in a large social network. New England Journal of Medicine, 2008(358.21):2249~2258

Voting Mobilization: A Facebook Study

- Voting mobilization [Bond et al, Nature'2012]
  - show a facebook msg. on voting day with faces of friends who voted
  - generate 340K additional votes due to this message, among 60M people tested

Bond R M, Fariss C J, Jones J J, Kramer A D I, Marlow C, Settle J E, and Fowler J H. A 61-million-person experiment in social influence and political mobilization. Nature, 2012(489):295~298

Influence Propagation Modeling and Optimizations

- How to model influence propagation in a social network?
  - Stochastic diffusion models
- How to optimize the influence propagation effect?
  - Influence maximization and its variants

- One core problem: Influence maximization
  - Find a small number of individuals in a network to generate a large influence
  - Applications in viral marketing, diffusion monitoring, rumor control, etc.

Influence Maximization in a Nutshell

Independent Cascade (IC) Model

- Each edge (u, v) has an *influence probability* $p(u, v)$
- Initially seed nodes in S_0 are activated
- At each step t, each node u activated at step $t-1$ activates its neighbor v independently with probability $p(u, v)$
- Influence spread $\sigma(S)$: expected number of activated nodes
- Other models: linear threshold (LT), triggering, general threshold, etc.

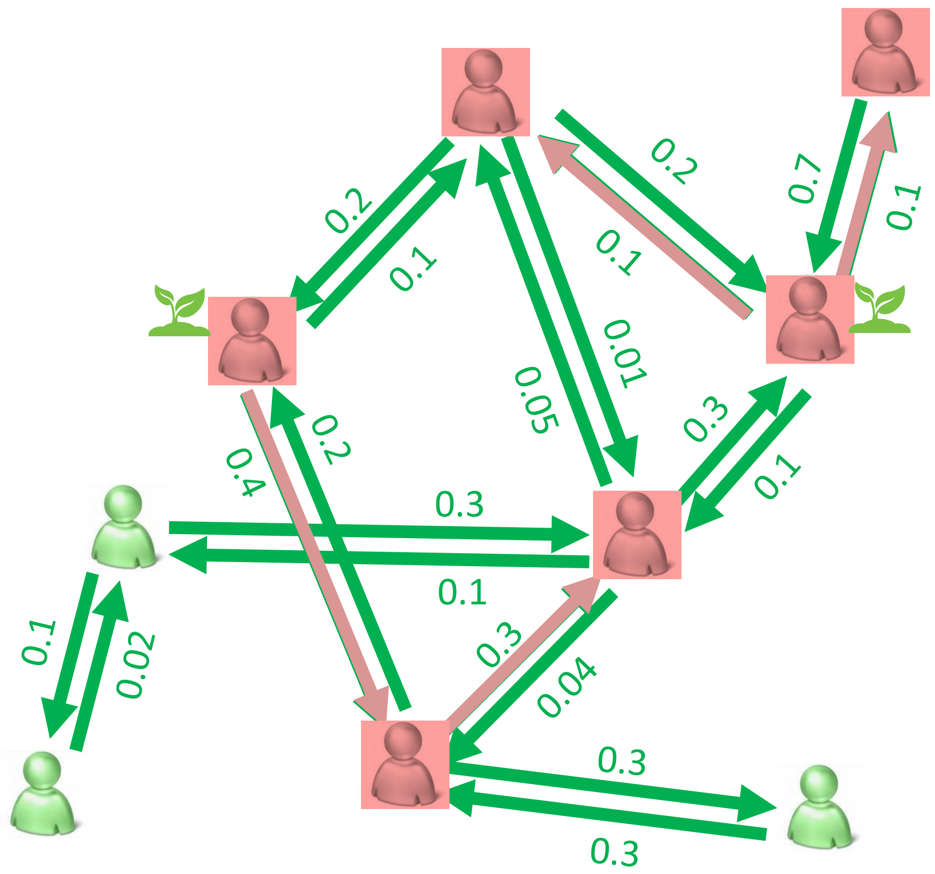

Influence Maximization

- Given a social network, a diffusion model with given parameters, and a number $\boldsymbol{k}$, find a seed set $\boldsymbol{S}$ of at most $\boldsymbol{k}$ nodes such that the influence spread of $\boldsymbol{S}$ is maximized.
- Based on *submodular function* maximization
- [Kempe, Kleinberg, and Tardos, KDD'2003]

Kempe D, Kleinberg J M, and Tardos É. Maximizing the spread of influence through a social network. KDD'2003

Submodular Set Functions

- Sumodularity of set functions $f: 2^V \to R$
  - for all $S \subseteq T \subseteq V$, all $v \in V \setminus T$,
    $$f(S \cup \{v\}) - f(S) \geq f(T \cup \{v\}) - f(T)$$
  - diminishing marginal return
  - an equivalent form: for all $S, T \subseteq V$
    $$f(S \cup T) + f(S \cap T) \leq f(S) + f(T)$$
- Monotonicity of set functions f: for all $S \subseteq T \subseteq V$,
  $$f(S) \leq f(T)$$

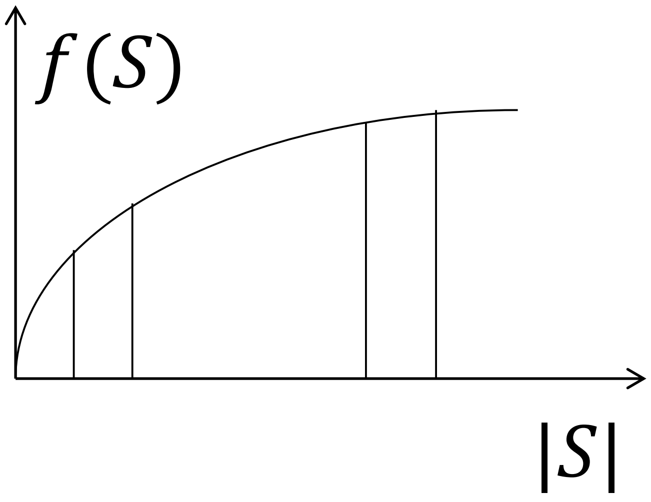

Submodularity of Influence Spread Function $\sigma(S)$

- Independent cascade model is equivalent to
  - sample live edges by edge probabilities
  - activate nodes reachable from S in the live-edge graph
- $\sigma(S) = \sum_L \Pr\{L\} \cdot |\Gamma(L,S)|$
  - $\Gamma(L,S)$: set of nodes reachable from S in live-edge graph L
  - $|\Gamma(L,S)|$ is a coverage function, easy to show it is submodular

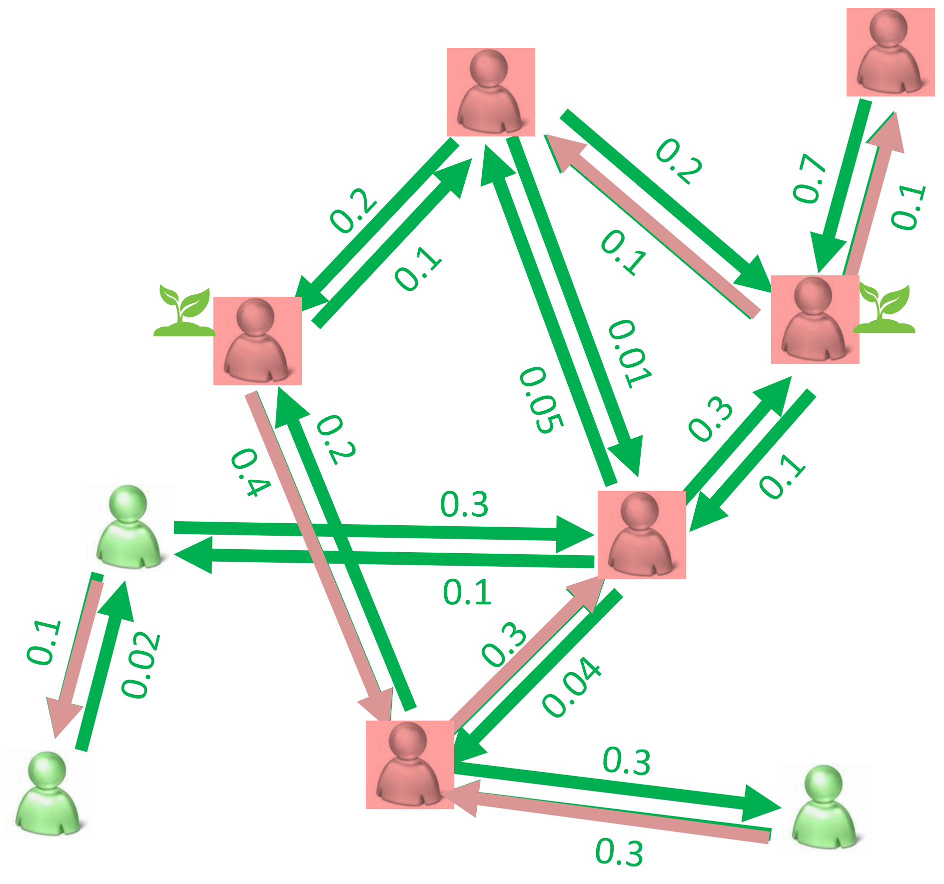

Greedy Algorithm for Submodular Function Maximization

initialize $S = \emptyset$;
for $i = 1$ to k do
    select $u = \mathrm{argmax}_{w \in V \setminus S}[f(S \cup \{w\}) - f(S))]$
    $S = S \cup \{u\}$
end for
output S

Property of the Greedy Algorithm

- Theorem: If the set function f is monotone and submodular with $f(\emptyset) \geq \mathbf{0}$, then the greedy algorithm achieves $(\mathbf{1} - \mathbf{1}/\boldsymbol{e})$ approximation ratio, that is, the solution $\boldsymbol{S}$ found by the greedy algorithm satisfies:

$$f(S) \geq \left(1 - \frac{1}{e}\right) \max_{S' \subseteq V, |S'| = k} f(S')$$

- [Nemhauser, Wolsey and Fisher, 1978]

Nemhauser G L, Wolsey L A, and Fisher M L. An analysis of approximations for maximizing submodular set functions. Mathematical Programming 1978

Challenges and Research Coverage

- Scalability challenge:
  - In IC (and LT) models, computing influence spread $\sigma(S)$ for any given S is #P-hard [Chen et al. KDD'2010, ICDM'2010], and Monte Carlo simulation is slow
  - Scalable influence maximization
- Adaptivity challenge:
  - Can we adapt to partial feedbacks? --- adaptive influence maximization
- Learning challenge:
  - How to learn the diffusion model?
  - How to use online feedback for optimization --- online influence maximization
- Complex model challenge:
  - Other variants of influence diffusion models --- competitive and complementary influence maximization, non-submodular influence maximization, etc.

Pushing the Limits of Optimization and Learning

- Influence maximization
  - Sitting at the boundary of feasibility
- Examples discussed in this talk
  - Adaptive influence maximization: new variants in adaptive maximization

	Adaptive submodular	Non-adaptive submodular
Independent feedback	prior studies	IC + myopic feedback
Dependent feedback	IC + full-adoption feedback	LT + myopic/full-adoption feedback

  - Online influence maximization --- general combinatorial multi-armed bandit framework
    - prior studies: linear rewards, exact offline oracle
    - Online IM: nonlinear rewards, approximation oracle, probabilistically triggered arms

Adaptive Influence Maximization

Adaptive Influence Maximization: Model

- Influence propagation: IC Model
- Seed selection: one-by-one, instead of a batch of k nodes
  - After selecting each seed node, obtain feedback on the propagation from the seed --- can be used to help subsequent seed selection
- Two feedback models
  - Full-adoption feedback: all downstream propagation from the selected seed, whether an edge passes through influence, whether a node is activated
  - Myopic feedback: only immediate propagation from the seed to its out-neighbors are given as the feedback

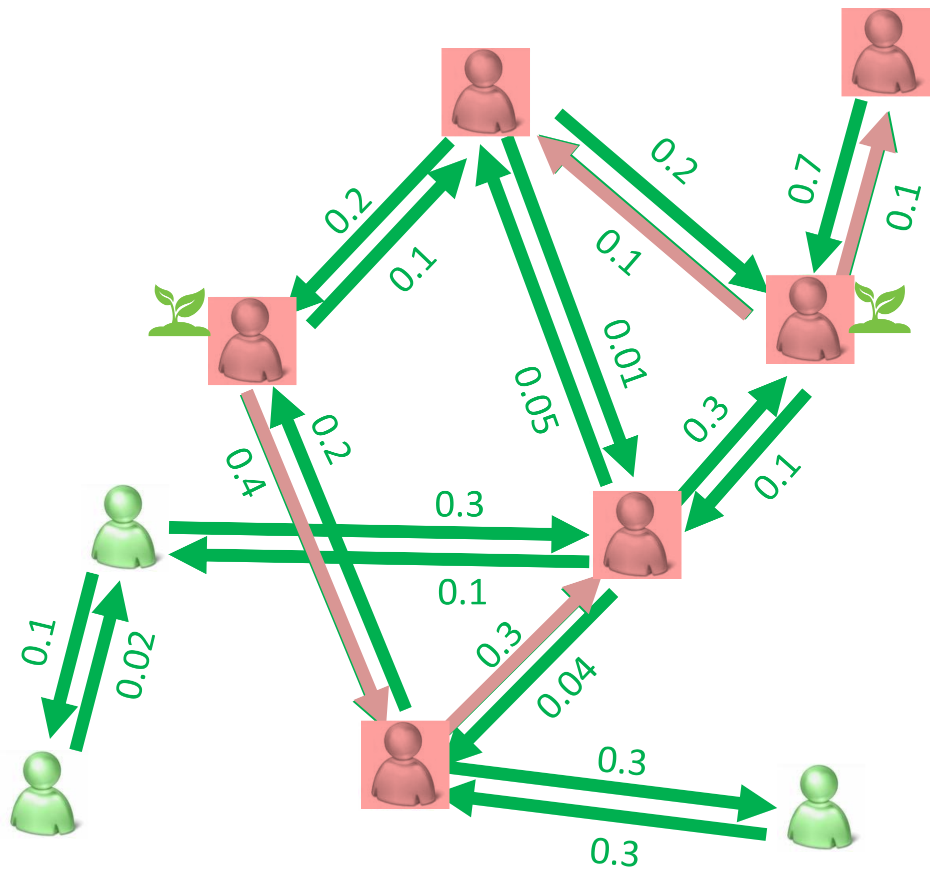

Full Adoption Feedback / Myopic Feedback

Adaptive Submodularity

- Realization ϕ: all randomness in a propagation (random live-edge graph)
- Partial realization ψ: feedback collected (partial propagation) from the currently selected seeds $\mathbf{dom}(\psi)$
- Adaptive Submodularity: a node u's marginal influence is higher on a smaller partial realization than on a larger partial realization
  – If $\psi \subseteq \psi'$, $\Delta(u|\psi') \leq \Delta(u|\psi)$
- Adaptive Monotonicity: a node u's marginal influence on any partial realization is nonnegative
  – $\Delta(u|\psi) \geq 0$, as long as ψ has non-zero probability to occur

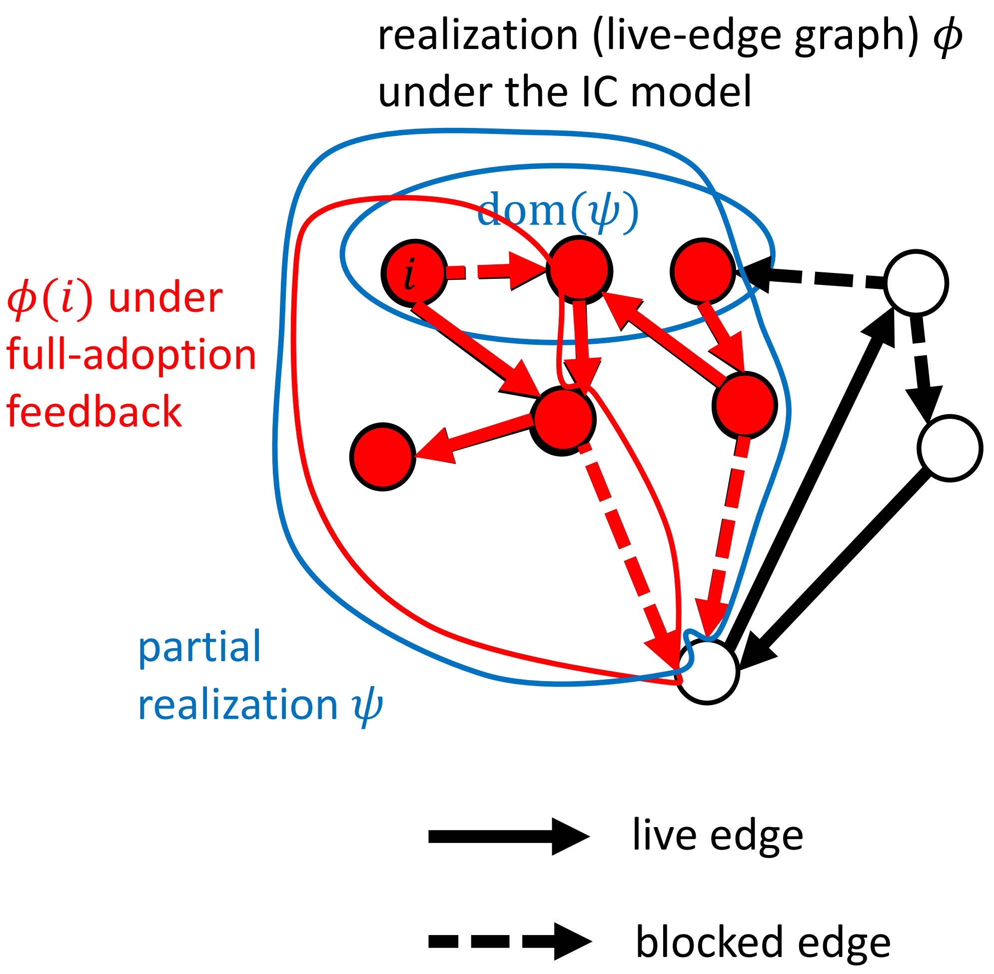

Important Results on Adaptive Submodularity

- Approximation ratio
  - Adaptive Greedy Algorithm
    - On every step, greedily select the next entry with the largest marginal influence: select $v = \mathrm{argmax}_u \Delta(u|\psi)$
  - If the model is adaptive monotone and adaptive submodular, adaptive greedy algorithm is a $1 - 1/e$ approximation of the adaptive optimal solution. [GK11]
- Adaptivity gap: supremum ratio of the adaptive optimal vs. non-adaptive optimal: $\sup_{G,k} \frac{OPT_A(G,k)}{OPT_N(G,k)}$
  - If the model is adaptive monotone and adaptive submodular, and the feedback are mutually independent, the adaptivity gap is at most $\frac{e}{e-1}$. [AN16]

Golovin D and Krause A. Adaptive submodularity: theory and applications in active learning and stochastic optimization. Journal of Artificial Intelligence Research, 2011
Asadpour A and Nazerzadeh H. Maximizing stochastic monotone submodular functions. Management Science, 2016

Adaptive Submodularity on Influence Maximization

	IC model	LT model
Full-adoption feedback	adaptive submodular	not adaptive submodular
Myopic feedback	not adaptive submodular	not adaptive submodular

- IC+myopic not adaptive submodular
  - Example to the right
  - [GK11] conjectures that adaptive greedy is still a constant approximation
  - We answer this conjecture affirmatively [NeurIPS'19]

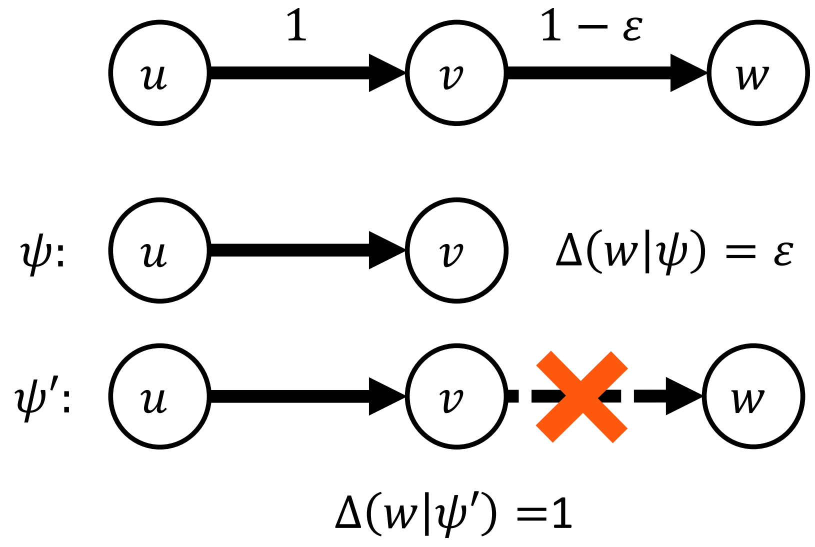

Adaptivity Gap on Influence Maximization

- Related to adaptive Submodularity and feedback independence

	Adaptive submodular	Non-adaptive submodular
Independent feedback	$\frac{e}{e-1}$ [AN16]	IC + myopic feedback $\left[\frac{e}{e-1}, 4\right]$ [PC19]
Dependent feedback	IC + full-adoption feedback Partial answers on specific graphs [CP19]: In-arborescences: $\left[\frac{e}{e-1}, \frac{2e}{e-1}\right]$ Out-arborescences: $\left[\frac{e}{e-1}, 2\right]$ Bipartite graphs: $\frac{e}{e-1}$	LT + myopic/full-adoption feedback ? Triggering + full-adoption: unbounded [CPST20] Triggering+myopic ?

Peng B and Chen W. Adaptive influence maximization with myopic feedback, NeurIPS'2019

Chen W and Peng B. On adaptivity gaps of influence maximization under the independent cascade model with full adoption feedback. ISAAC'2019

Chen W, Peng B, Schoenebeck G, and Tao B. Adaptive greedy versus non-adaptive greedy for influence maximization. AAAI'2020

Implications on IC + Myopic Feedback

- Adaptive greedy is $\frac{1}{4} \cdot \left(1 - \frac{1}{e}\right)$ approximation of the adaptive optimal solution
  - Adaptive greedy is $\left(1 - \frac{1}{e}\right)$ approximation of the non-adaptive optimal
  - Non-adaptive optimal is ¼ of the adaptive optimal
- Answers the open conjecture by Golovin and Krause (2011)
- First study on adaptive maximization with a non-adaptive submodular model

Idea on the Analysis, IC + Myopic, Gap ≤ 4

- Decision tree representation of adaptive policy $\boldsymbol{\pi}$
- Random walk non-adaptive policy $\boldsymbol{\mathcal{W}(\pi)}$: Select $\mathbf{dom}(\boldsymbol{\psi_\ell})$ with probability $\boldsymbol{p_\ell}$
- Fictitious hybrid policy $\bar{\boldsymbol{\pi}}$ and aggregate adaptive influence spread $\bar{\boldsymbol{\sigma}}(\bar{\boldsymbol{\pi}})$
- Show: $\boldsymbol{\sigma}(\boldsymbol{\pi}) \leq \bar{\boldsymbol{\sigma}}(\bar{\boldsymbol{\pi}}) \leq 4\boldsymbol{\sigma}(\boldsymbol{\mathcal{W}}(\boldsymbol{\pi}))$

Decision tree of policy π

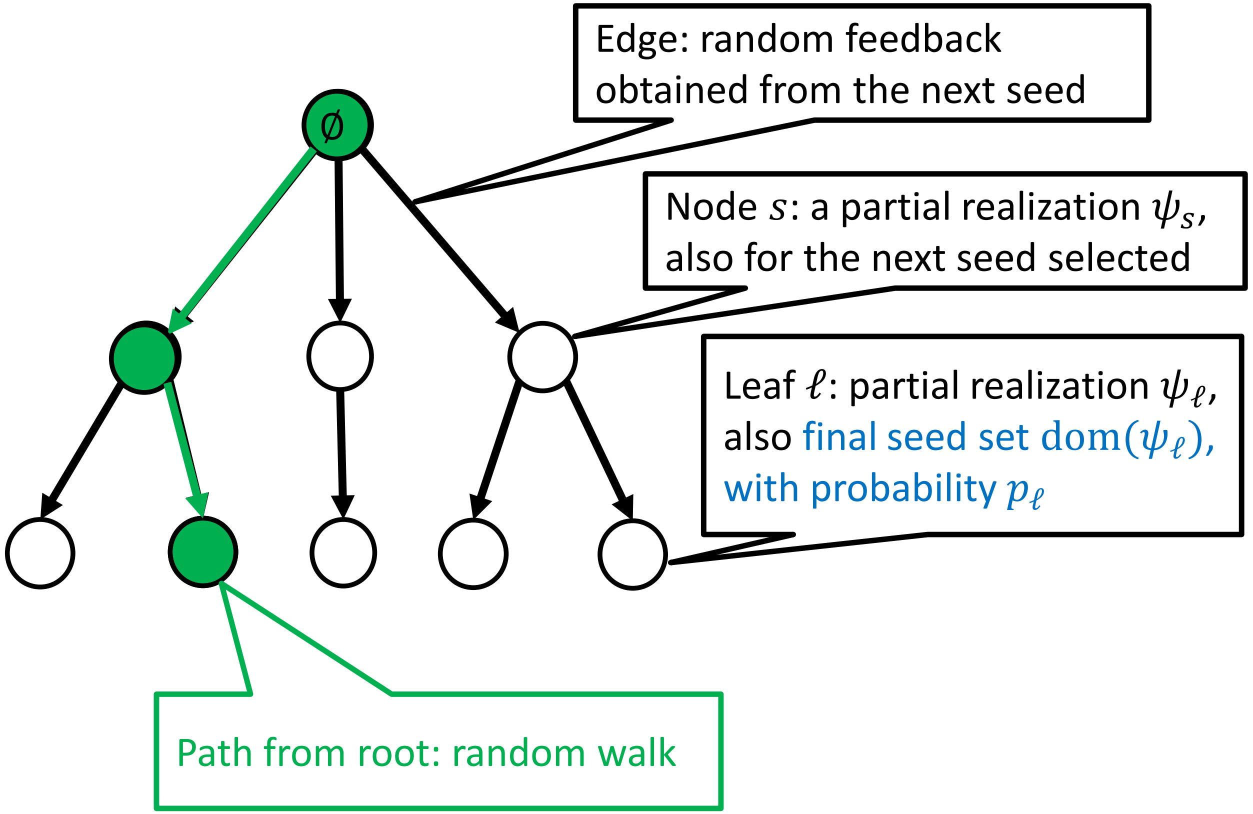

Fictitious Hybrid Policy $\bar{\pi}$ and Aggregate Adaptive Influence Spread $\bar{\sigma}(\bar{\pi})$

- Work simultaneously on three independent realizations Φ_1, Φ_2, Φ_3
- $\bar{\pi}$ selects seeds adaptively exactly like π working on Φ_1
- But for each selected seed u, it has three independent chances to activate its out-neighbors, according to Φ_1, Φ_2, Φ_3
- The expected number of activated nodes is the aggregate adaptive influence spread $\bar{\sigma}(\bar{\pi})$
- Obviously, $\sigma(\pi) \leq \bar{\sigma}(\bar{\pi})$

Connecting Aggregate Adaptive Spread $\bar{\sigma}(\bar{\pi})$ with Non-Adaptive Spread $\sigma(\mathcal{W}(\pi))$

- t-th aggregate influence spread $\sigma^t(S)$ and t-th aggregate adaptive influence spread $\sigma^t(\pi)$, $t = 1, 2, 3$
  - each seed gets t independent chances of activating its out-neighbors
  - $\bar{\sigma}(\bar{\pi}) = \sigma^3(\pi)$, $\sigma(\mathcal{W}(\pi)) = \sigma^1(\mathcal{W}(\pi))$
- Represent $\sigma^t(\mathcal{W}(\pi))$ and $\sigma^t(\pi)$ by non-adaptive marginal gains $\Delta_{f^t}(u|\mathrm{dom}(\psi_s))$ and adaptive marginal gains $\Delta_{f^t}(u|\psi_s)$, respectively
  - telescoping series on node s along a path in the decision tree
- $\Delta_{f^3}(u|\psi_s) \le 2\Delta_{f^2}(u|\mathrm{dom}(\psi_s)) \Rightarrow \sigma^3(\pi) \le 2\sigma^2(\mathcal{W}(\pi))$
  - key lemma, crucially relying on (a) feedback independence and (b) (non-adaptive) submodularity of influence utility function on live-edge graphs
- $\sigma^2(\mathcal{W}(\pi)) \le 2\sigma(\mathcal{W}(\pi))$

Greedy Adaptivity Gap

- Motivation:
  - optimal solutions cannot be reached
  - Practical algorithms are mostly greedy-based
- Greedy Adaptivity Gap: ratio between adaptive greedy vs. non-adaptive greedy
- Results:

	IC model	**LT model**	**Triggering model (more general)**
Full-adoption	tight lower bound: $1-1/e$	tight lower bound: $1-1/e$	Upper bound: unbounded
Myopic	tight lower bound: $1-1/e$	tight lower bound: $1-1/e$	

Chen W, Peng B, Schoenebeck G, and Tao B. Adaptive greedy versus non-adaptive greedy for influence maximization. AAAI'2020

Many Open Problems

- Adaptivity gap:

	IC model	LT model	Triggering model (more general)
Full-adoption	result on general graphs? tighter result for special graphs?	?	unbounded
Myopic	$\left[\frac{e}{e-1}, 4\right]$, tight result?	?	?

- Greedy adaptivity gap:

	IC model	LT model	Triggering model (more general)
Full-adoption	upper bound?	upper bound?	lower: $1 - 1/e$ upper: unbounded
Myopic	upper bound $\leq \frac{4e}{e-1}$ tight upper bound?	upper bound?	upper bound?

- Better adaptive algorithms than greedy?

Online Influence Maximization

Microsoft

Online Influence Maximization

- Edge influence probabilities are unknown, need to be learned
- Multiple rounds of online influence maximization. In each round,
  - select $\boldsymbol{k}$ seeds to influence the network
  - observe the diffusion paths and results
  - collect the reward --- the number of nodes activated
  - use the observed feedback to update learning statistics, which is used for seed selection in later rounds
- Falls into the online learning (multi-armed bandit) framework

Multi-Armed Bandit Problem

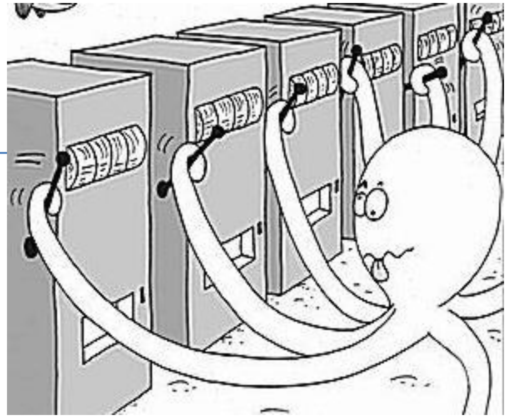

- There are m arms (machines)
- Arm i has an unknown reward distribution on $[0,1]$ with unknown mean μ_i
  - best arm $\mu^* = \max \mu_i$
- In each round, the player selects one arm to play and observes the reward
- Performance metric: Regret:
  - Regret after playing T rounds $= T\mu^* - \mathbb{E}[\sum_{t=1}^{T} R_t(i_t^A)]$
- Objective: minimize regret in T rounds
- Balancing exploration-exploitation tradeoff
  - exploration: try new arms
  - exploitation: keep playing the best arm so far
- Wide applications: Any scenario requiring selecting best choice from online feedback
  - online recommendations, advertising, wireless channel selection, social networks, A/B testing

Classical MAB Algorithm: UCB1

1: for each arm i: $\hat{\mu}_i = 1$ (empirical mean), $T_i = 0$ (number of observation)

2: for $t = 1, 2, 3, \ldots$ do

3: for each arm i: $\rho_i = \sqrt{\frac{3 \ln t}{2T_i}}$ (confidence radius)

4: for each arm i: $\bar{\mu}_i = \min\{\hat{\mu}_i + \rho_i, 1\}$ (upper confidence bound, UCB)

5: $j = \text{argmax}_i \, \bar{\mu}_i$

6: play arm j, observe its reward $X_{j,t}$

7: update $\hat{\mu}_j = (\hat{\mu}_j \cdot T_j + X_{j,t})/(T_j + 1)$; $T_j = T_j + 1$

6: end-for

For exploration (ρ_i)

For exploitation ($\hat{\mu}_i$)

Guarantee of the UCB1 Algorithm

- Finite-horizon regret:
  - distribution dependent: $O\left(\sum_{\Delta_i>0}\frac{1}{\Delta_i}\ln T\right)$, $\Delta_i=\mu^*-\mu_i$
  - distribution independent: $O(\sqrt{mT\ln T})$

- [Auer, Cesa-Bianchi, and Fischer, 2002]

Auer P, Cesa-Bianchi N, and Fischer P. Finite-time analysis of the multiarmed bandit problem. Machine Learning Journal, 2002(47.2-3):235~256

Challenges Applying UCB1 to Online IM

- exponential number of seed sets
  - cannot treat each seed set as an arm
- non-linear reward functions
- offline problem is already NP-hard
- probabilistically triggering new arms in a play

Extending the MAB Framework

- Extend MAB to combinatorial MAB framework with probabilistically triggered arms (CMAB-T)
  - Model: In each round one action/super-arm is played, which triggers a set of base arms (triggering may be probabilistic)
  - precisely characterize the bounded smoothness condition required to solve CMAB-T
  - propose the CUCB algorithm based on an offline approximation oracle
  - distribution-dependent and distribution-independent regret analysis
  - applicable to a large class of combinatorial online learning problems
- [Chen et al JMLR'2016, Wang and Chen, NIPS'2017]

Chen W, Wang Y, Yuan Y, and Wang Q. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. Journal of Machine Learning Research, 2016

Wang Q and Chen W. Improving regret bounds for combinatorial semi-bandits with probabilistically triggered arms and its applications. NIPS'2017

CUCB Algorithm

for each arm i: $\hat{\mu}_i = 1$ (empirical mean), $T_i = 0$ (number of observation)

for $t = 1, 2, 3, \ldots$ do

for each arm i: $\rho_i = \sqrt{\frac{3 \ln t}{2T_i}}$ (confidence radius)

for each arm i: $\bar{\mu}_i = \min\{\hat{\mu}_i + \rho_i, 1\}$ (upper confidence bound, UCB)

$S = \mathrm{OfflineOracle}(\bar{\mu}_1, \ldots, \bar{\mu}_m)$

play action/super-arm S, observe triggered arm outcomes $\{X_{j,t}\}$

for each observed j: update $\hat{\mu}_j = (\hat{\mu}_j \cdot T_j + X_{j,t})/(T_j + 1)$; $T_j = T_j + 1$

end-for

Regret Bounds

- $O\left(\sum_i \frac{1}{\Delta_{\min}^i} B_1^2 K \ln T\right)$ distribution-dependent regret
  - i: base arm index
  - B_1: one-norm bounded-smoothness constant
  - K: maximum number of arms any action can trigger
  - T: time horizon, total number of rounds
  - $\Delta_{\min}^i$: minimum gap between α fraction of the optimal reward and the reward of any action that could trigger arm i (α is the offline approximation ratio)
- $O\left(B_1 \sqrt{mKT \ln T}\right)$ distribution-independent regret
- For influence maximization, B_1 is the largest number of nodes any node can reach
- Main technical contribution: (a) characterizing the exact conditions (b) more sophisticated techniques in the analysis

Open Problems in Online Influence Maximization

- More efficient algorithms specific to influence maximization?
- Regret lower bound?
- Online IM with nonstationary distribution?

Conclusion and Future Work

- Influence maximization is a rich ground for studying many optimization and learning tasks
  - Right at the boundary of feasibility --- pushing the tasks to new limits
  - Other directions beyond adaptive and online influence maximization:
    - Scalable algorithms (well studied)
    - Learnability of influence functions (some studies)
    - Optimization from samples (initial study)
    - Non-submodular influence maximization (some studies)
    - Influence maximization + game theory (some initial studies)

Reference Resources

- Search “Wei Chen Microsoft”
  - Monograph: “Information and Influence Propagation in Social Networks”, Morgan & Claypool, 2013
  - my papers and talk slides
  - My upcoming book (in Chinese): 大数据网络传播模型和算法 (Network Diffusion Models and Algorithms for Big Data)

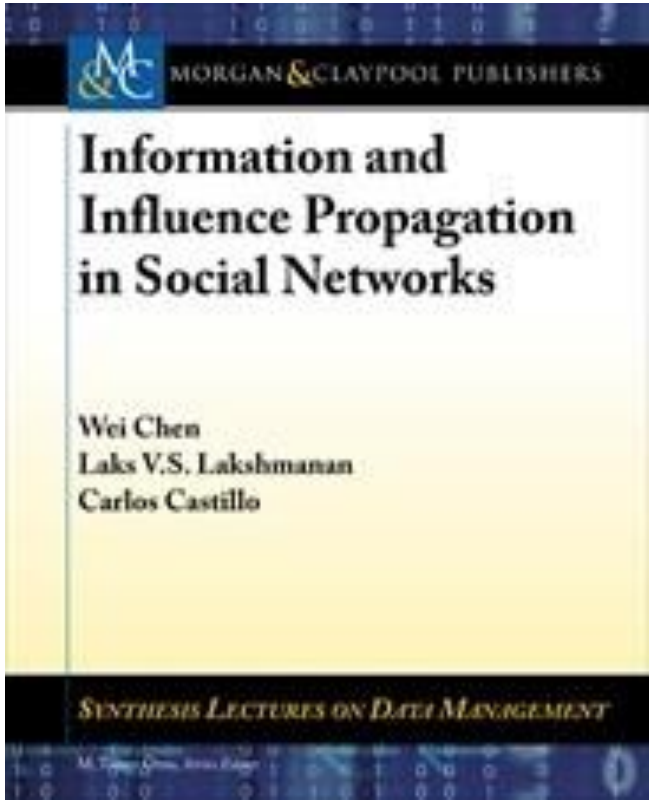

Thanks!