Computational Social Influence

Wei Chen
陈卫
Microsoft Research Asia

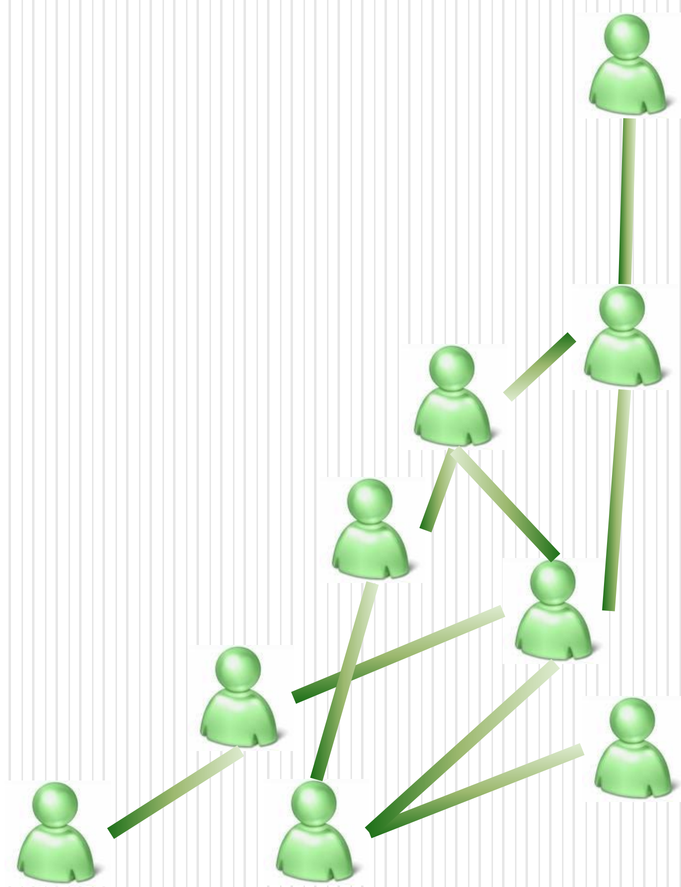

SocInf Workshop, IJCAI'2015, July 27, 2015

Social influence

- **Social influence** occurs when one's emotions, opinions, or behaviors are affected by others.

WIKIPEDIA
The Free Encyclopedia

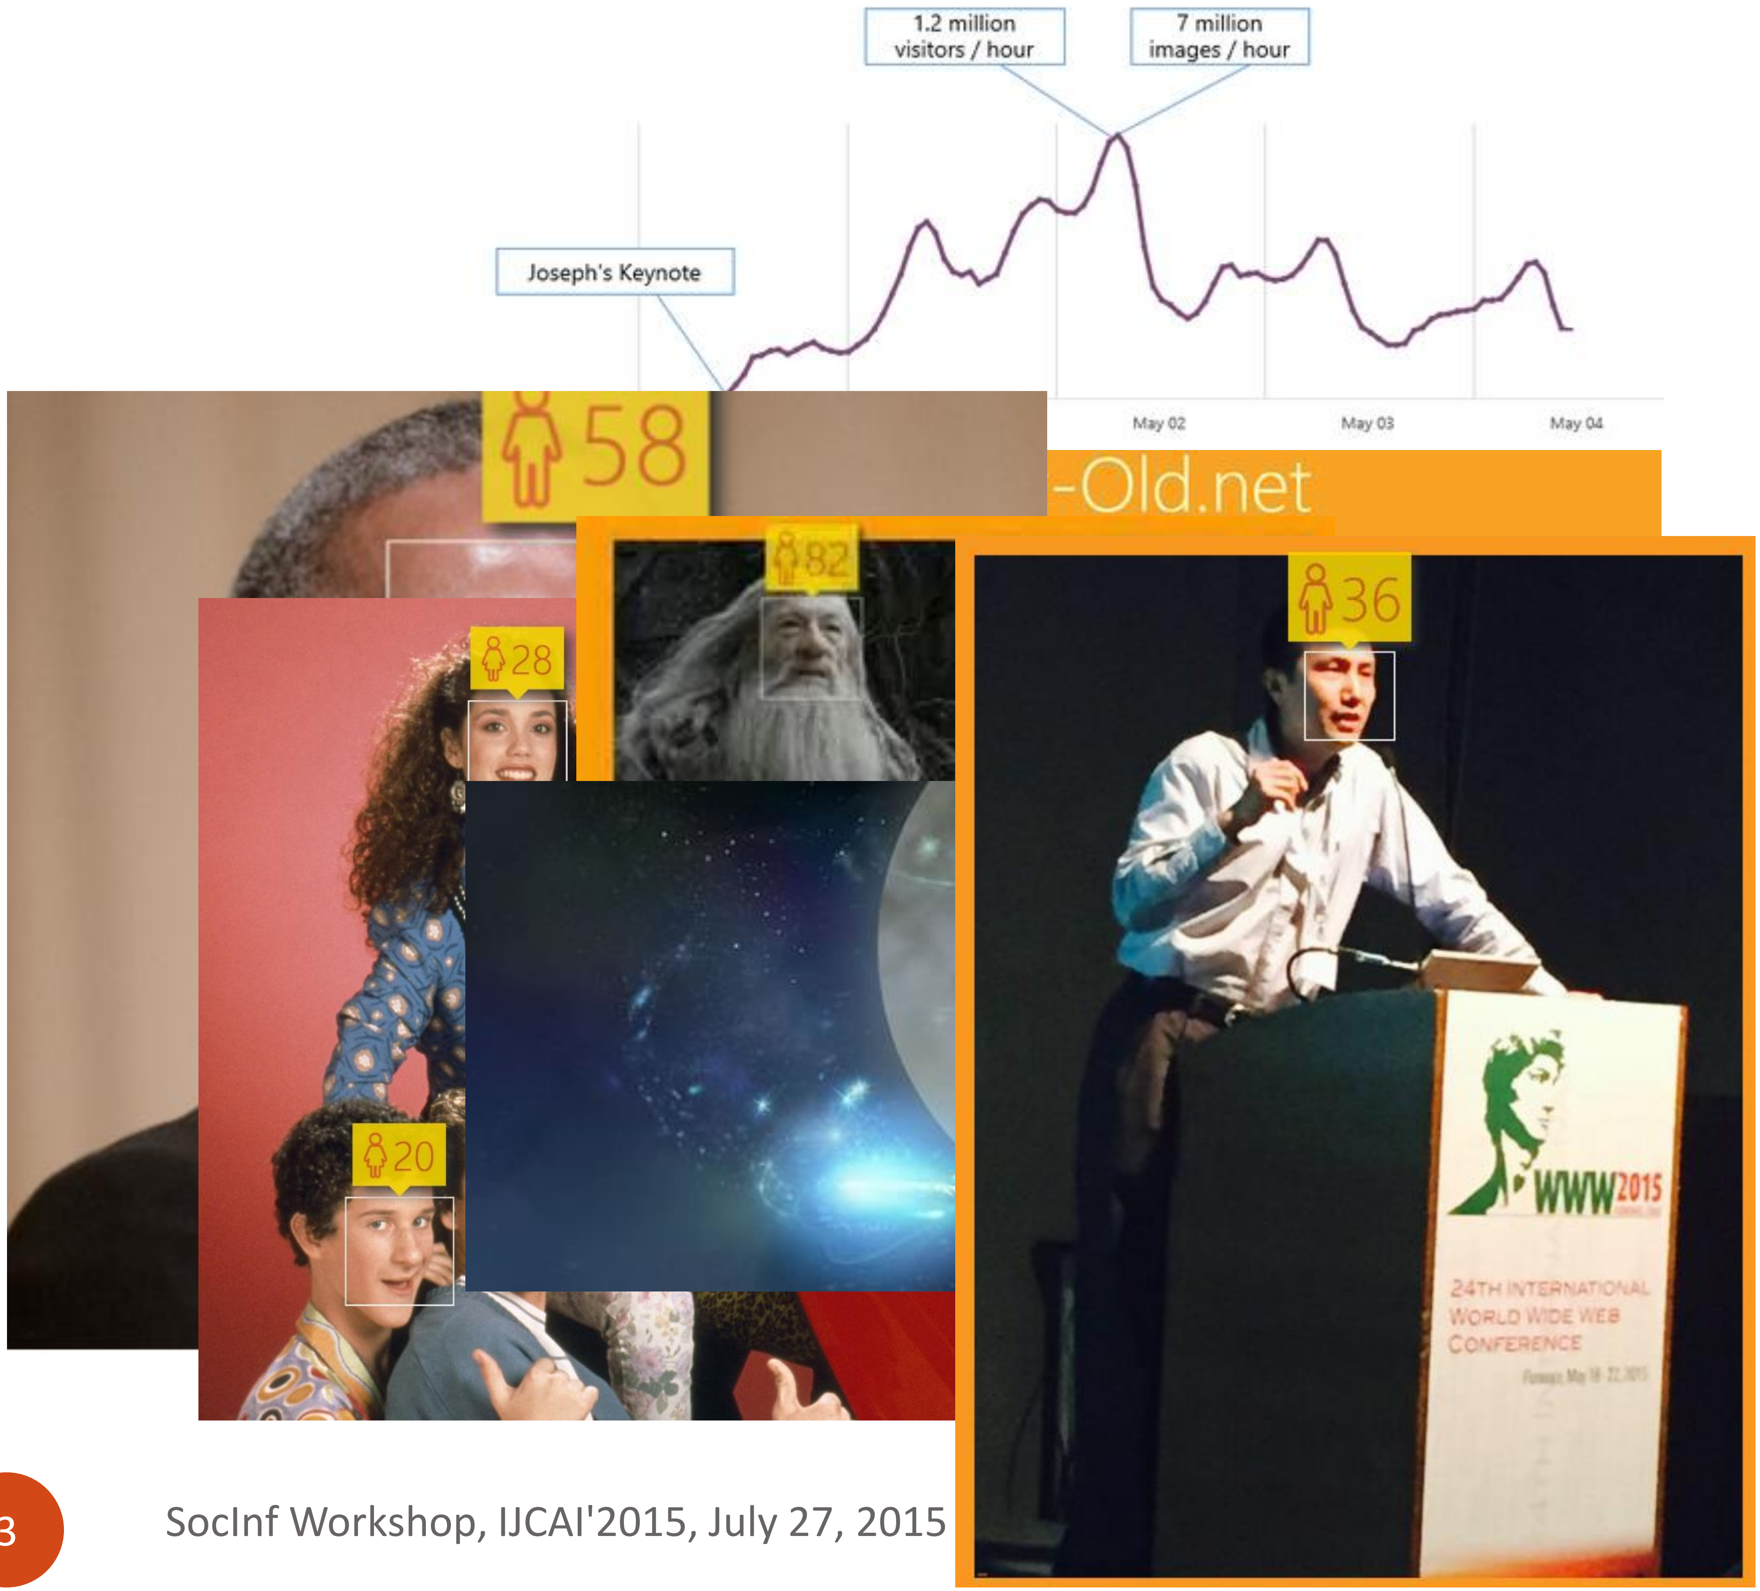
1.2 million visitors / hour
7 million images / hour
Joseph's Keynote
May 02
May 03
May 04
58
-Old.net
82
28
20
36
WWW2015
24TH INTERNATIONAL WORLD WIDE WEB CONFERENCE


MA5: 3857.54
MA10: 3789.76
MA20: 4125.02
MA30: 4432.35
5300
4900
4500
4100
3700
3300
5178.19
3373.54
2015/05
2015/06
2015/07

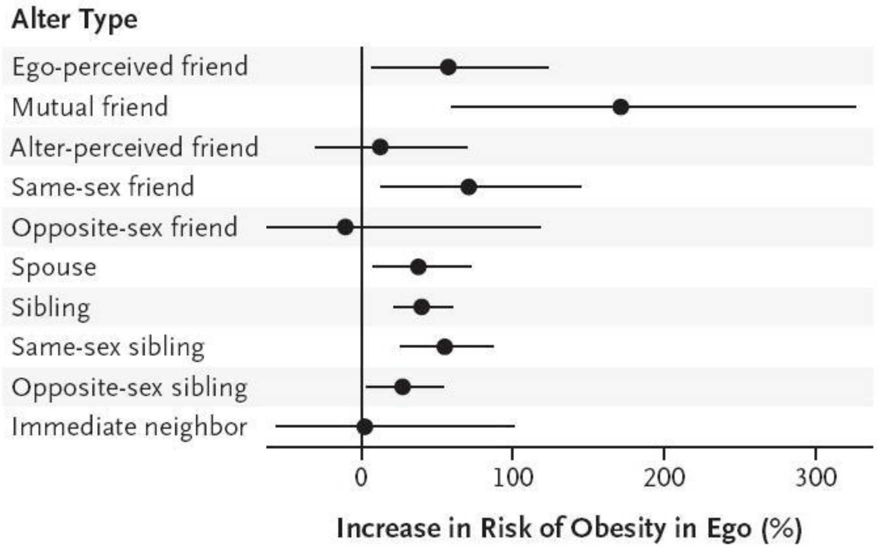

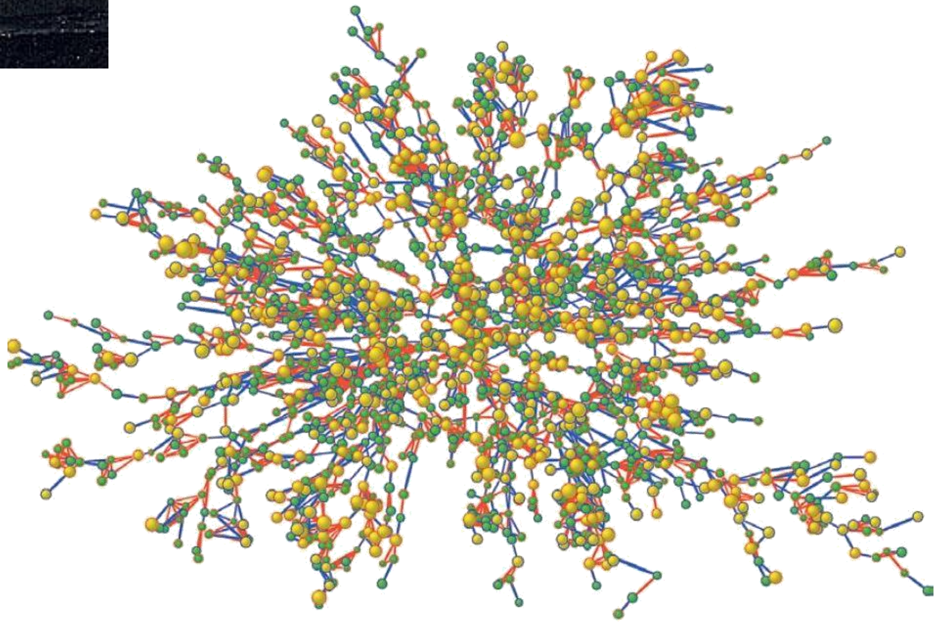

[Christakis and Fowler, NEJM'07,08]

Booming of online social networks

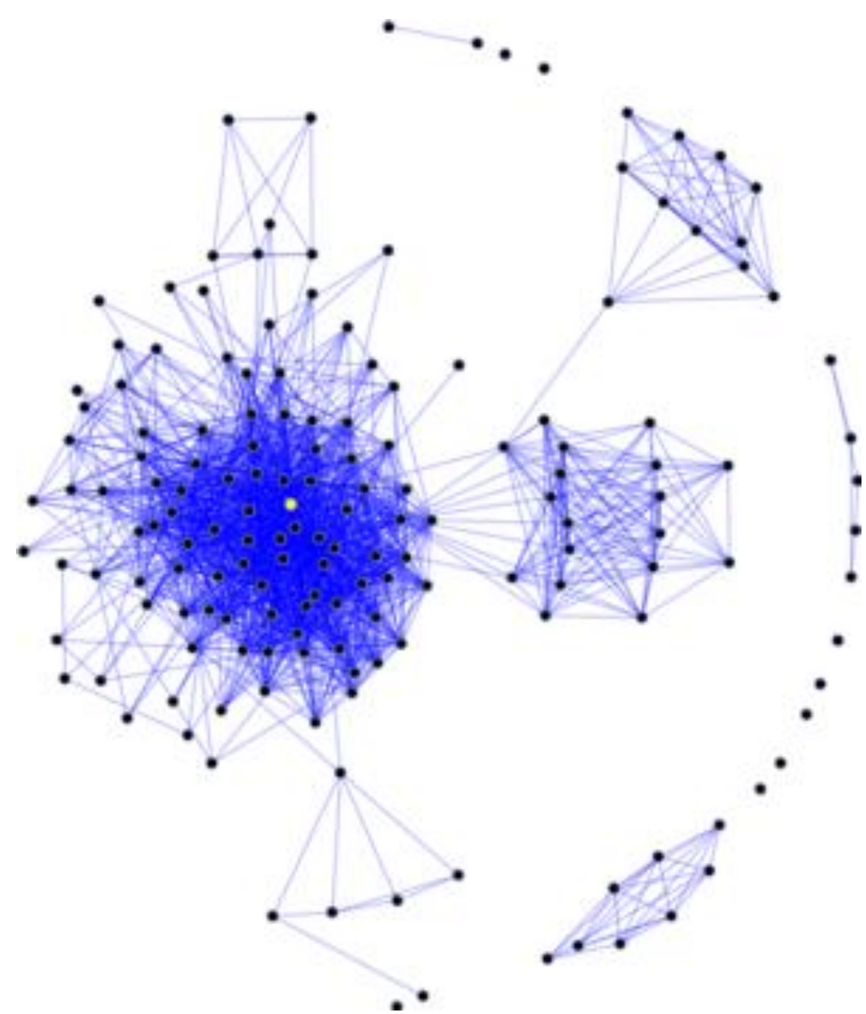

Opportunities on online social influence research and applications

- massive data set, real time, dynamic, open
- help understand social interactions, influence propagation patterns in grand scale
- help identifying influencers
- facilitate decision making in health care, business, politics, and economics

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

Three pillars of computational social influence

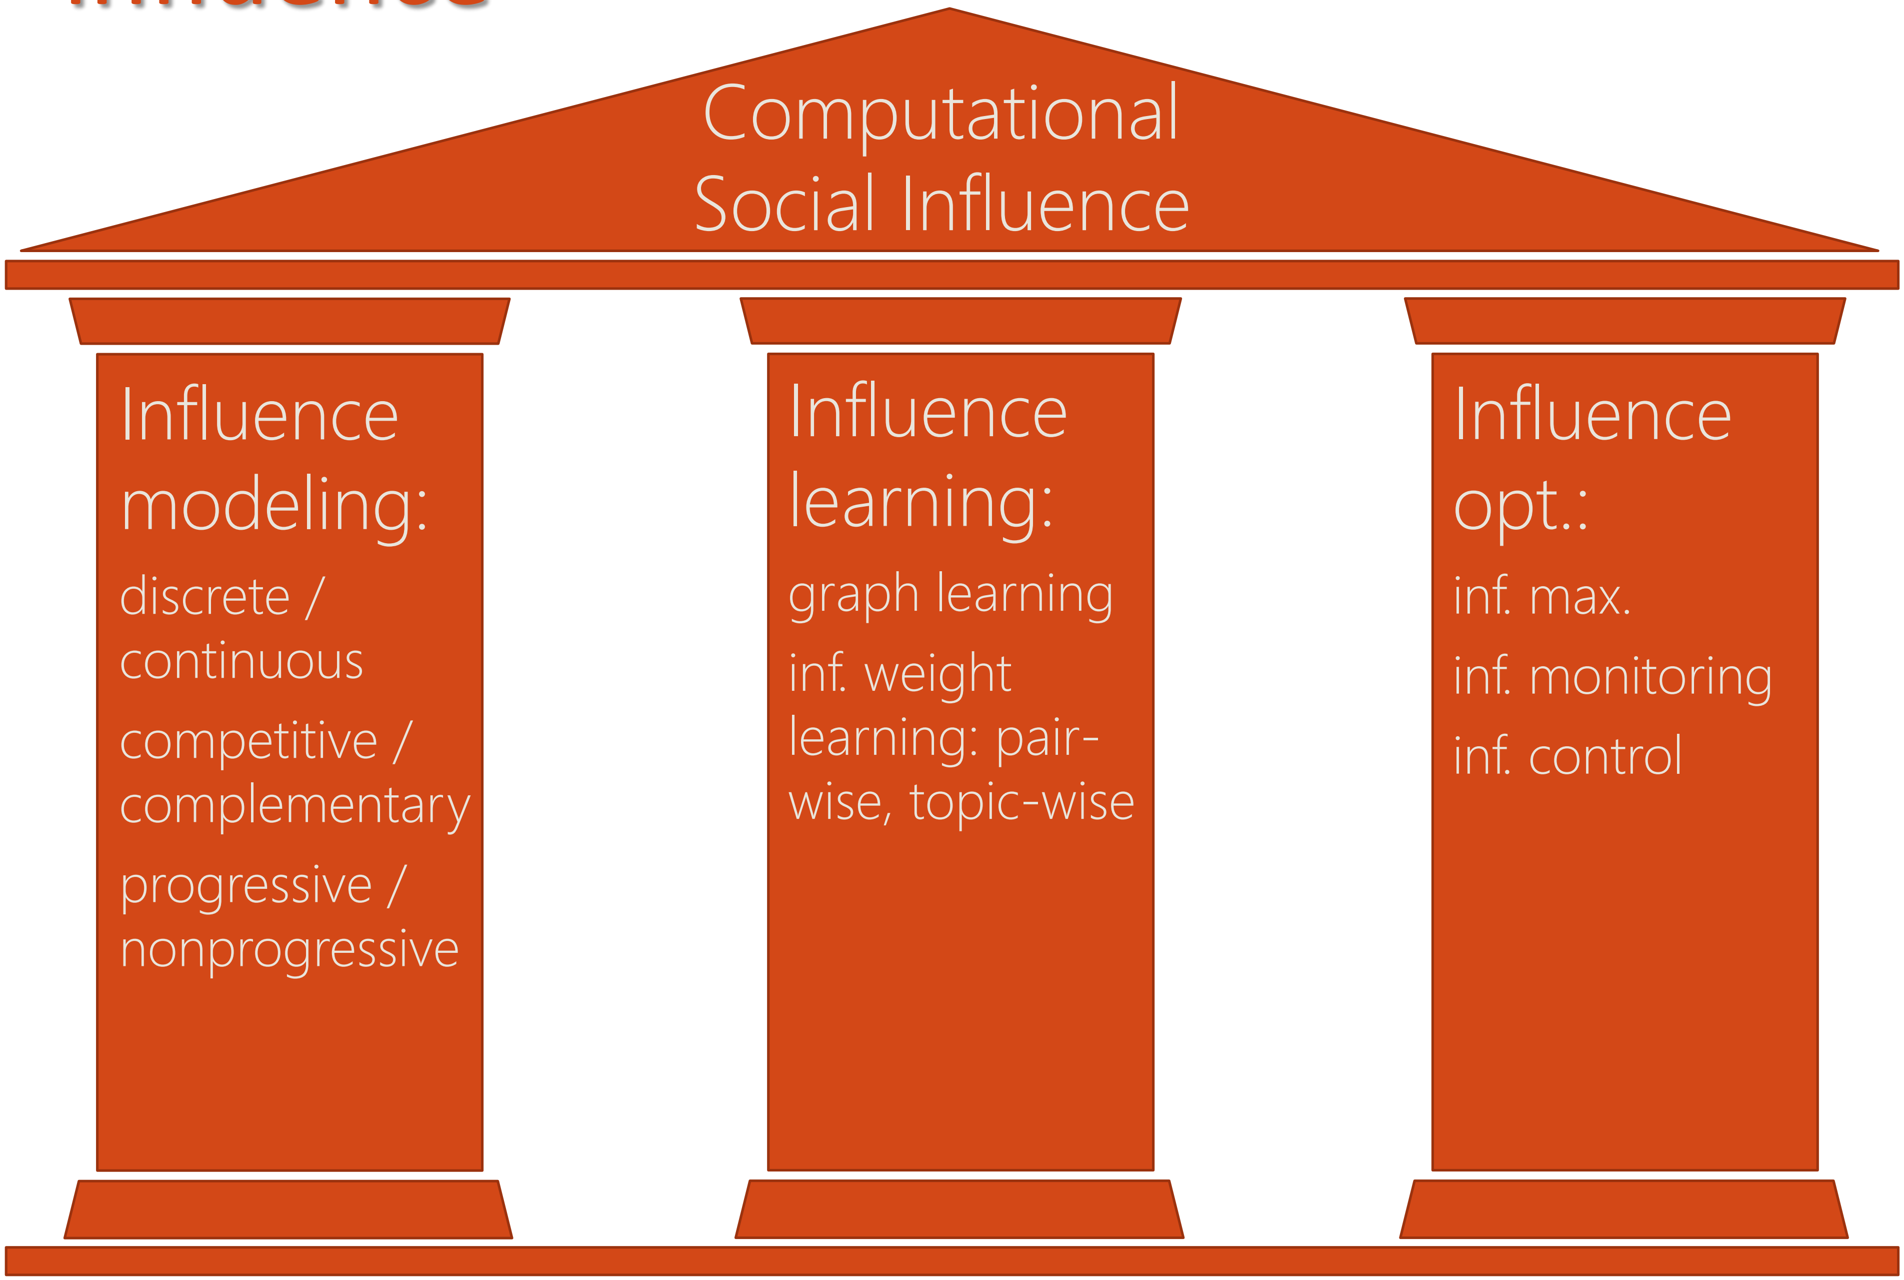

Influence modeling

- Discrete-time models:
  - independent cascade (IC), linear threshold (LT), general cascade models [KKT‘03]
  - topic-aware IC/LT models [BBM’12]
- Continuous-time models [GBS‘11]
- Competitive diffusion models
  - competitive IC [BAA‘11], competitive LT [HSCJ‘12], etc.
- Competitive & complementary diffusion model [LCL‘15]
- Others, epidemic models (SIS/SIR/SIRS…), voter model variants

Influence optimization

- Scalable inf. max.
  - Greedy approximation [KKT'03, LKGFVG'07, CWY'09, BBCL'14, TXS'14, TSX'15]
  - Fast heuristics [CWY'09, CWW'10, CYZ'10, GLL'11, JHC'12, CSHZC'13]
- Multi-item inf. max. [BAA'11, SCLWSZL'11, HSCJ'12, LBGL'13, LCL'15]
- Non-submodular inf. max. [GL'13, YHLC'13, ZCSWZ'14, CLLR'15]
- Topology change for inf. max. [TPTEFC'10,KDS'14]
- Inf. max with online learning [CWY'13, LMMCS'15]
- many others …

Influence learning

- Based on user action / adoption traces
- Learning the diffusion graph [GLK’10]
- Learning (the graph and) the parameters
  - frequentist method [GBL’10]
  - maximum likelihood [SNK’08]
  - MLE via convex optimization [ML’10,GBS’11,NS’12]

Rest of this talk

- Quick review of influence model and maximization
- Amphibious influence maximization
- Comparative influence diffusion --- from competition to complementarity

Quick Review of Influence Model and Maximization

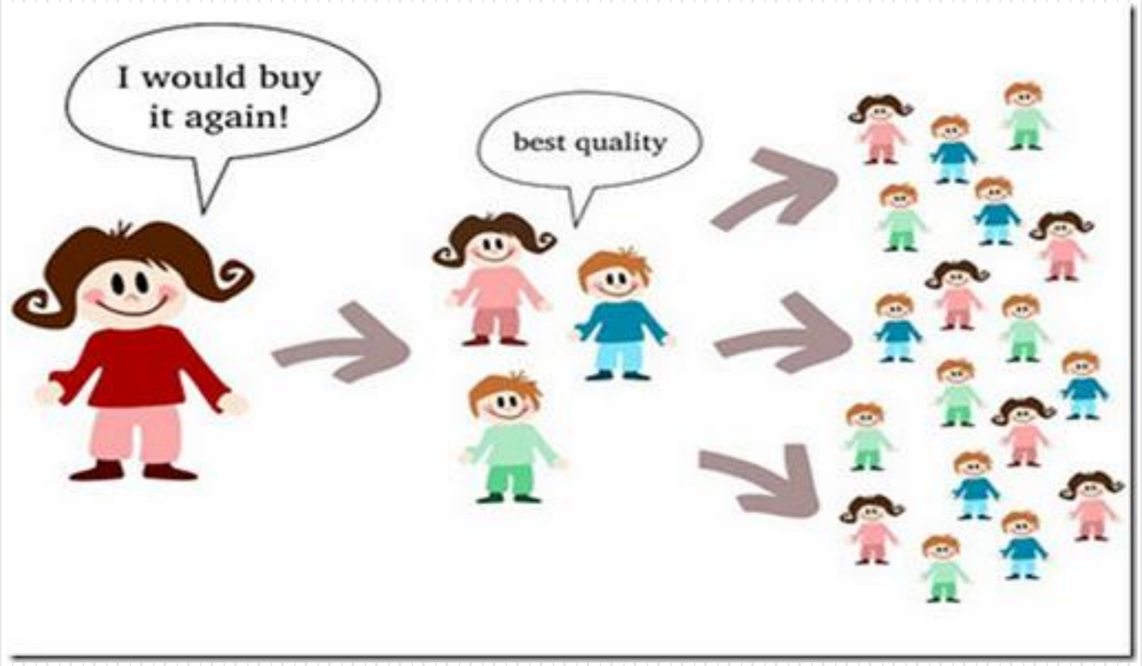

Example: viral marketing in social networks

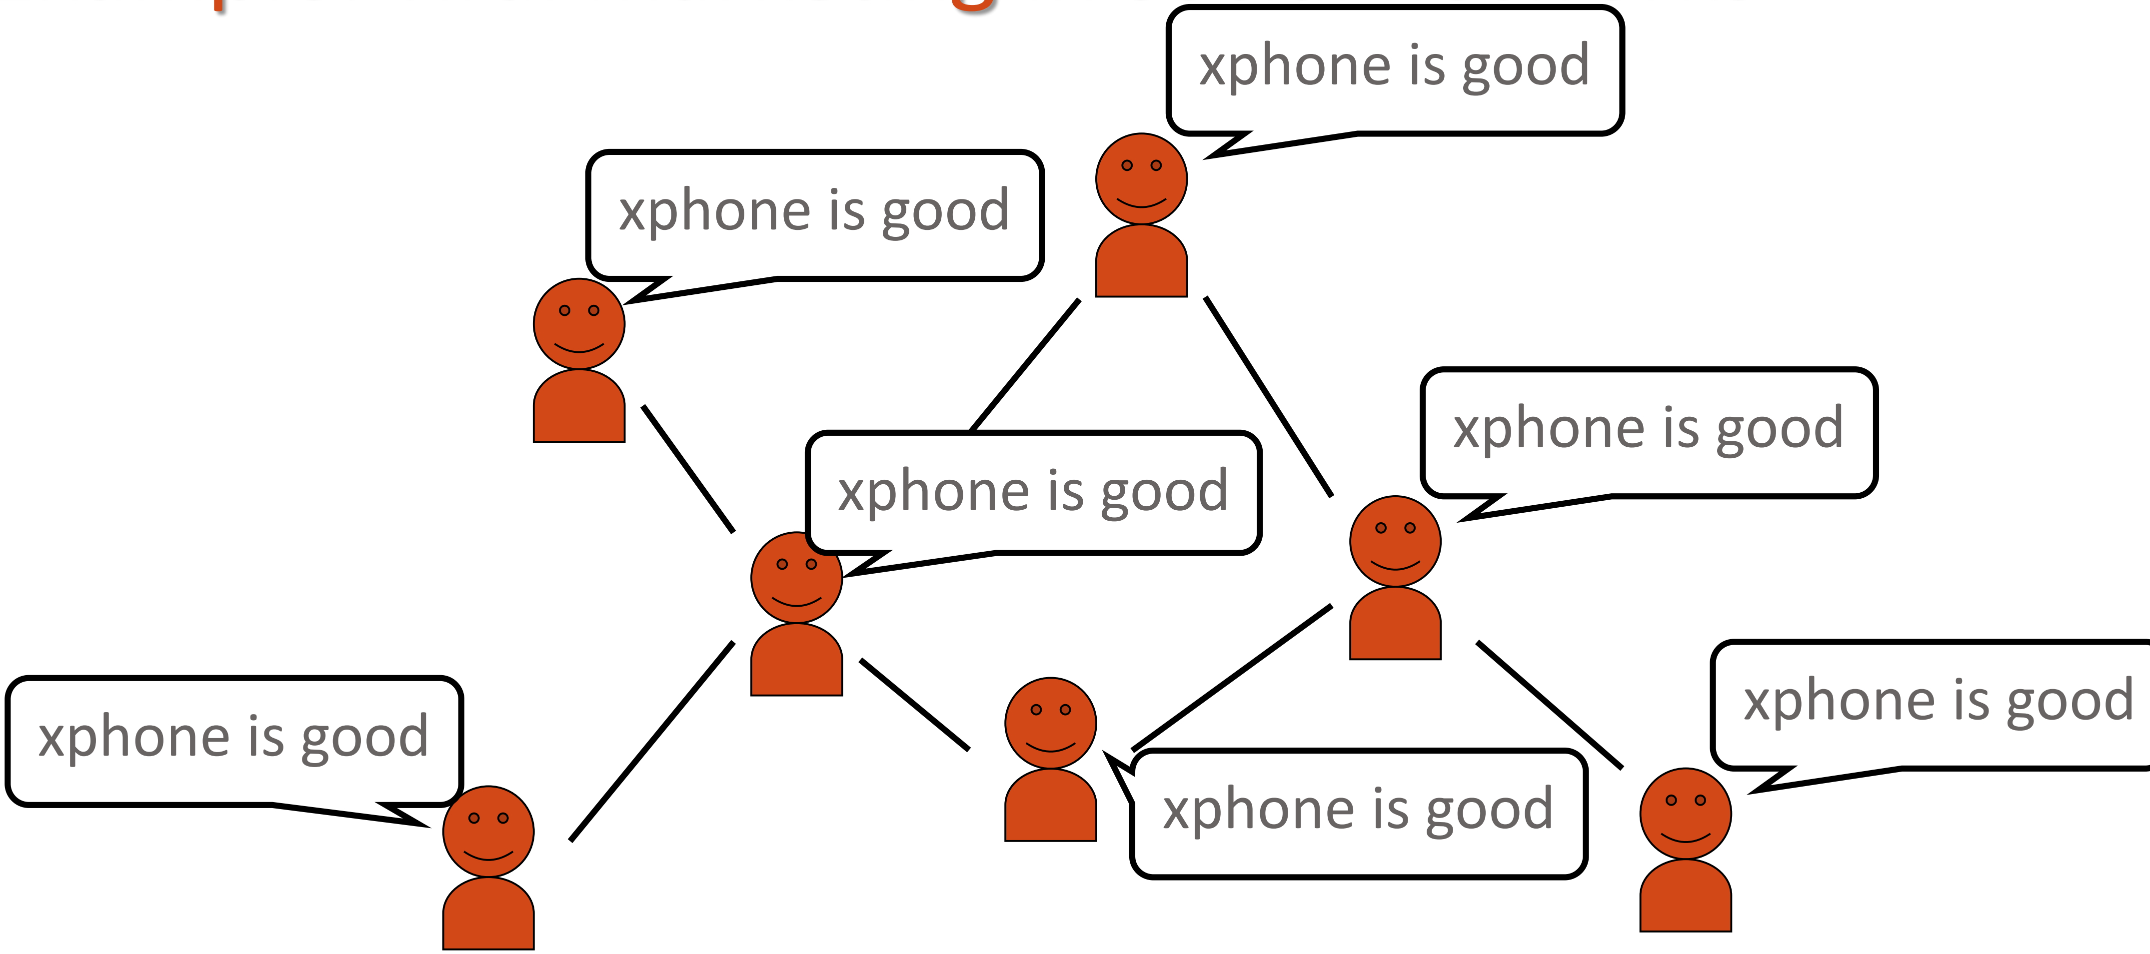

- Viral (word-of-mouth) marketing is believed to be a promising advertising strategy.
- Increasing popularity of online social networks may enable large scale viral marketing

Influence diffusion model

- Directed graph $G = (V, E)$
  - V: set of nodes, representing users
  - E: set of directed edges, representing influence relationships
- Influence probabilities on edges
  - $p(u, v)$: the probability that u activates v
- Independent cascade model
  - Initially nodes in a seed set S are activated
  - At step t, each node u activated at step $t-1$ has one chance to activate each of its out-going neighbor v, with success probability $p(u, v)$
  - influence spread $\sigma(S)$: expected number of active nodes when S is the seed set

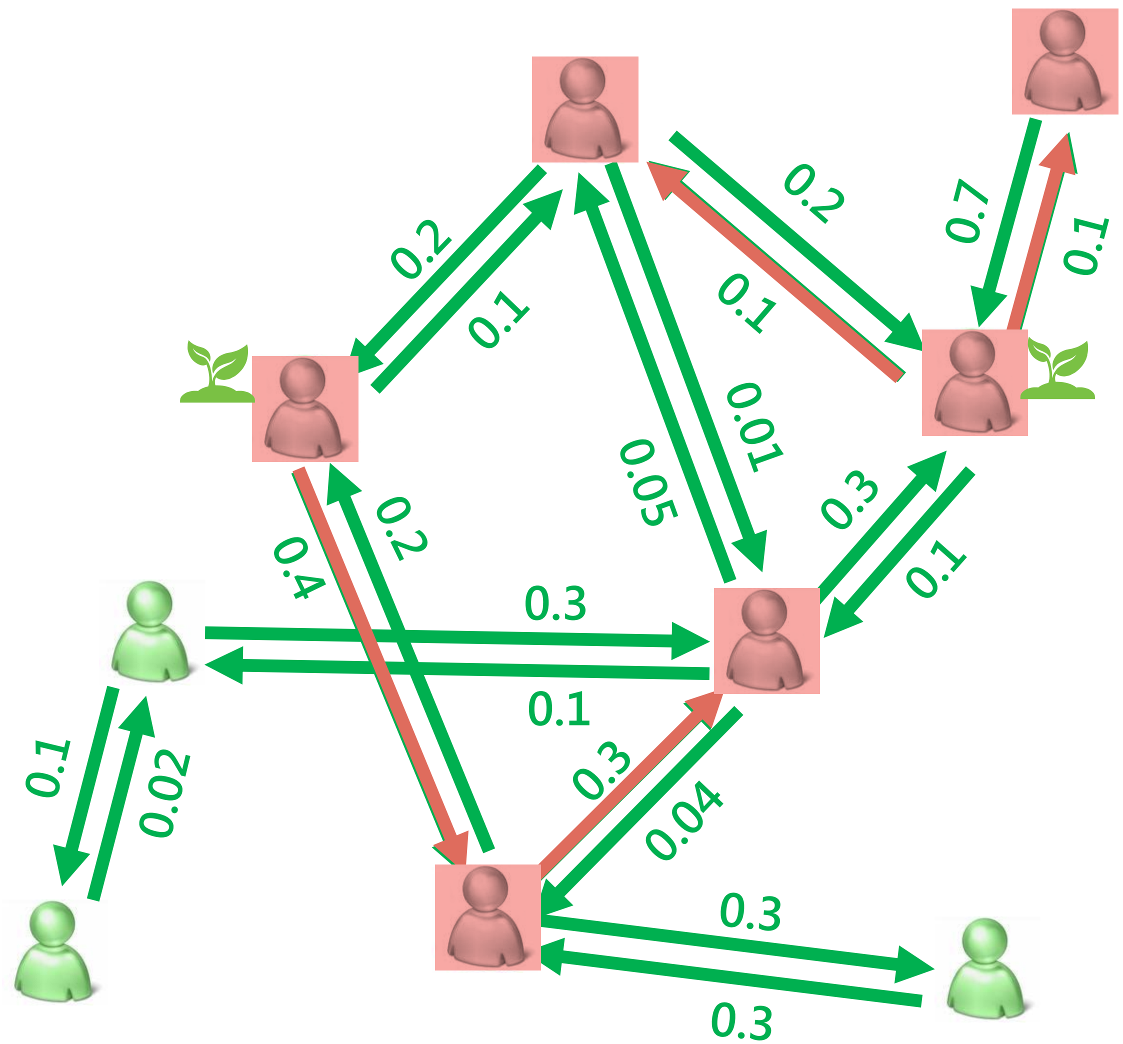

Influence maximization

- Given a social network, a diffusion model with given parameters, and a number k, find a seed set S of at most k nodes such that the influence spread of S, $\sigma(S)$, is maximized.
- Many possible variants
  - non-uniform cost, minimize seed set size, etc.
- NP hard

Greedy approximation framework

- influence spread $\sigma(S)$ is submodular:
  - for all $S \subseteq T \subseteq V$, all $v \in V \setminus T$,
    $\sigma(S \cup \{v\}) - \sigma(S) \geq \sigma(T \cup \{v\}) - \sigma(T)$
  - diminishing marginal return
- Submodular function maximization
  - Greedy algorithm: iteratively finding the next seed with the largest marginal influence spread
  - $1 - 1/e$ approximation

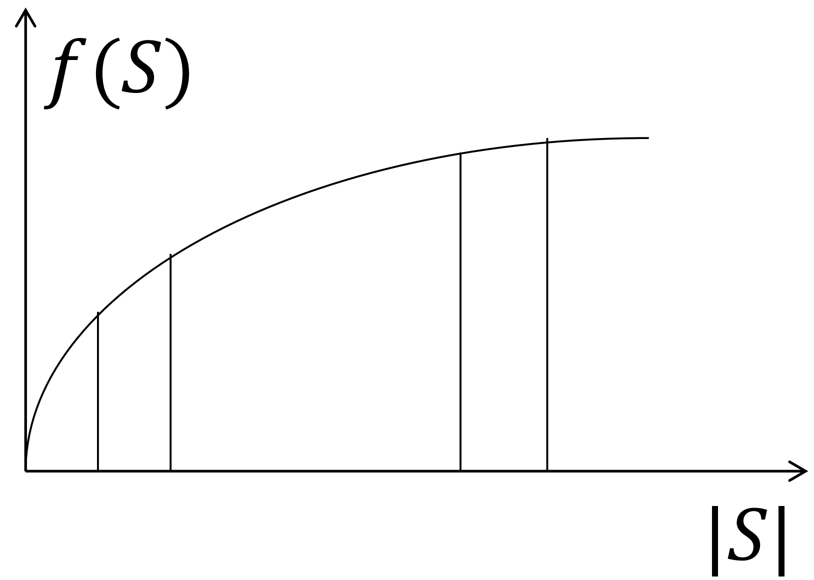

Fast MIA heuristic

- Pure greedy algorithm is very slow
- MIA uses local tree structure to replace general influence computation
- 1000 fold speedup with similar seed quality
- new algorithms available

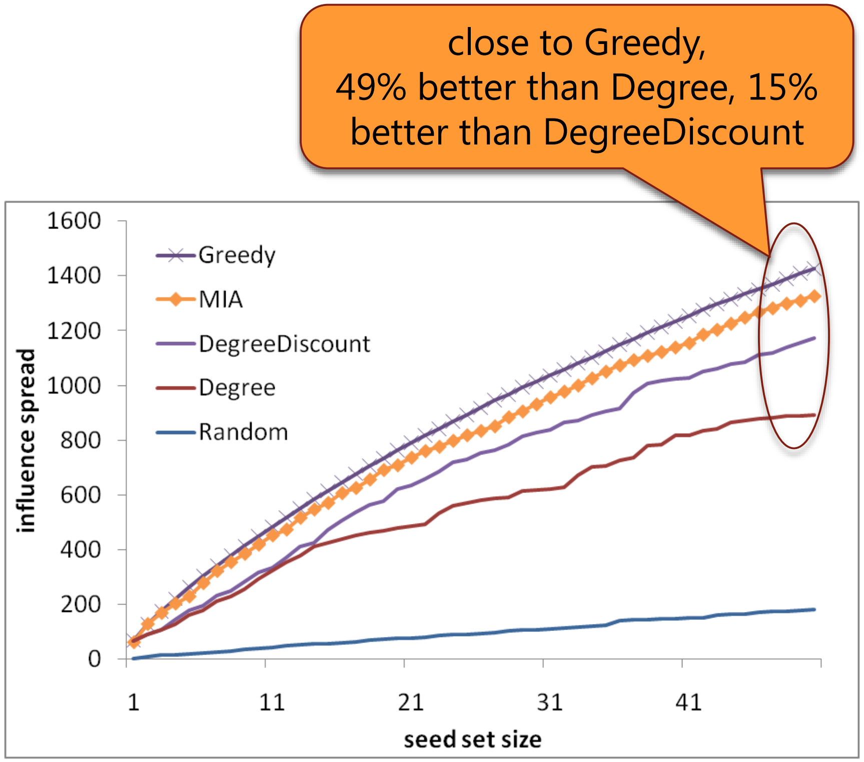

Influence spread vs. seed set size

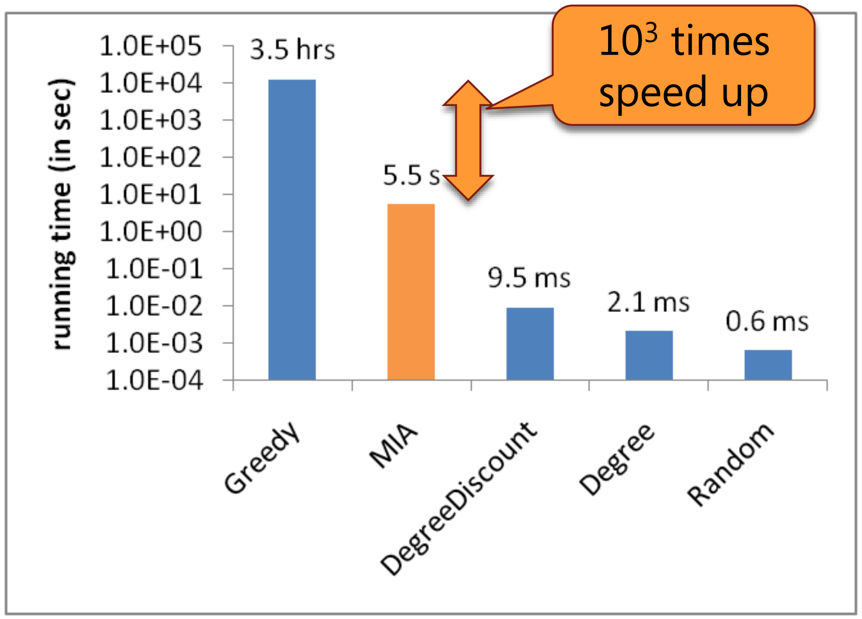

running time

Amphibious Influence Maximization

with Fu Li, Tian Lin, and Aviad Rubinstein

Traditional vs. viral marketing

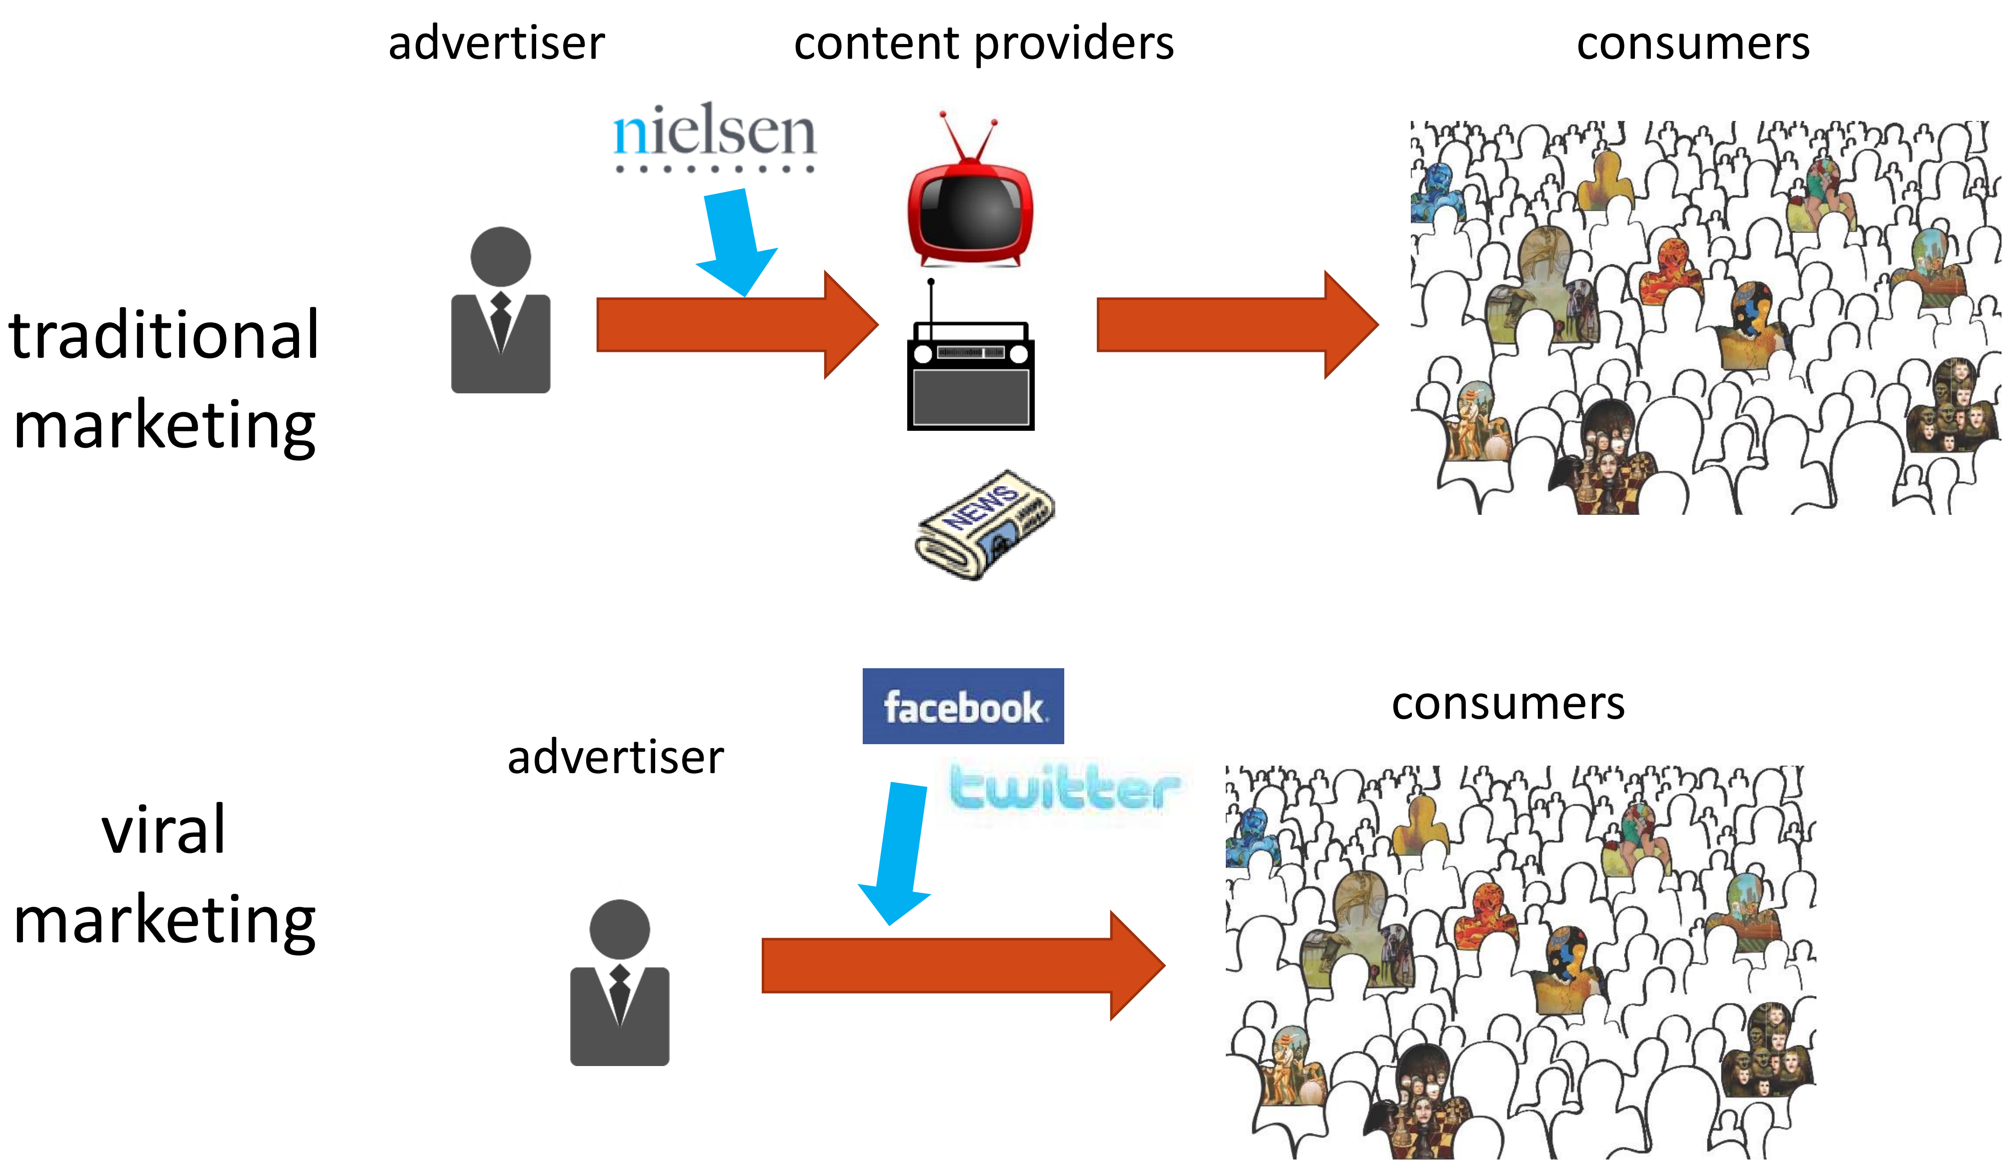

Amphibious marketing

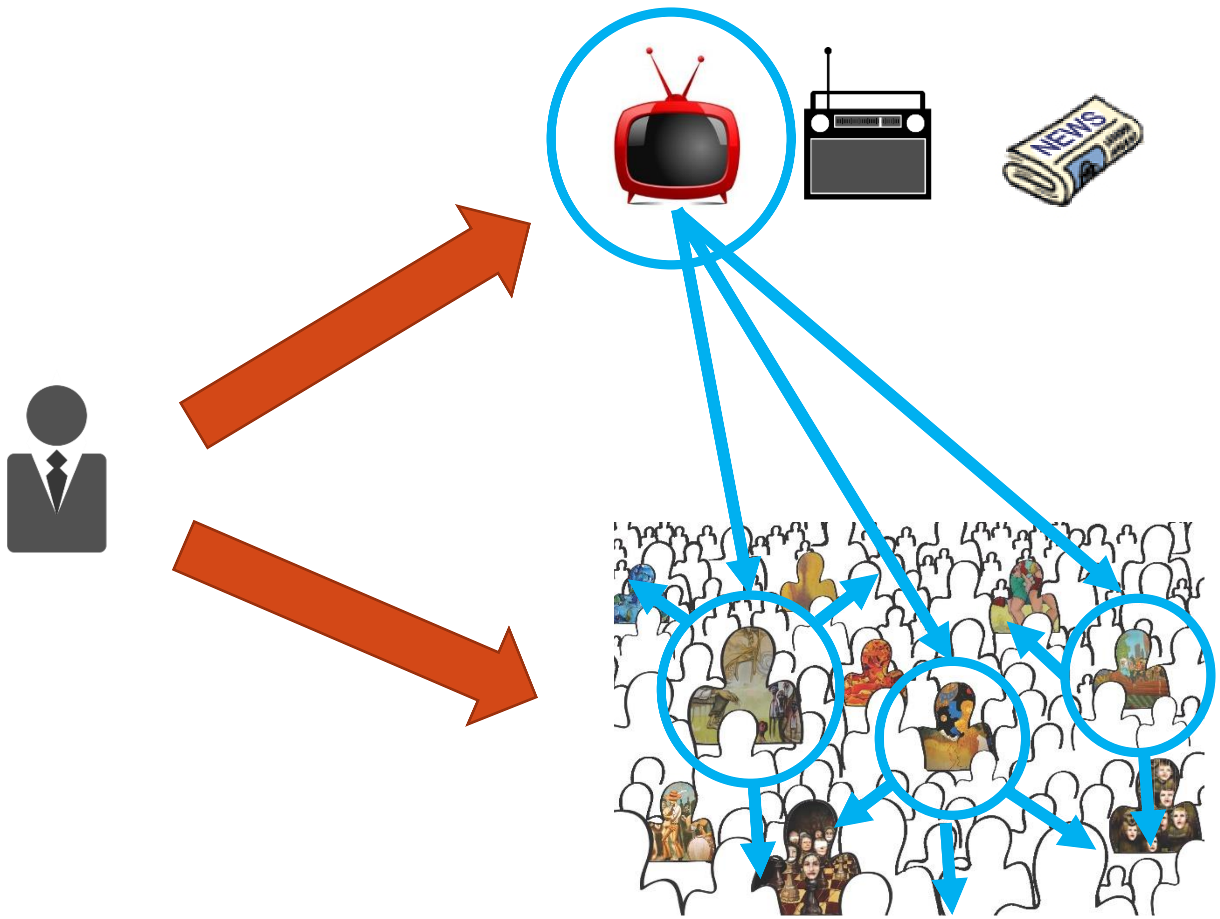

Amphibious influence maximization

- Given:
  - U – content providers
  - V – consumers
  - $G = (V, P)$ – social network graph with influence probability matrix P
  - $B = (U, V, M)$ – provider to consumer influence graph, with bi-adjacency matrix M
- Diffusion model:
  - From selected seed providers $X \subseteq U$, to selected seed consumers $Y \subseteq V$, then to the rest of social network
  - Same as IC model, with X as seeds, and remove all edges from U pointing to $V \setminus Y$
  - Influence spread $\sigma(X, Y)$: expected number of active nodes in V
- Goal
  - Find b_1 seed providers X and b_2 seed consumers Y to maximize $\sigma(X, Y)$

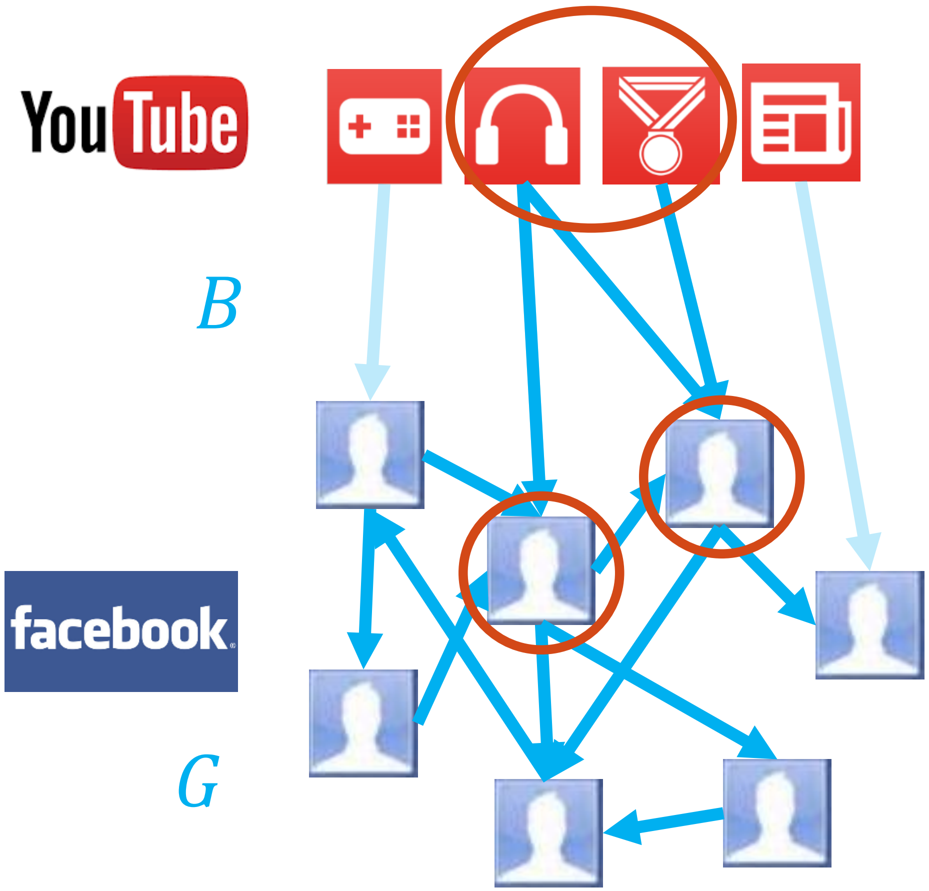

Main results

- Difficulty:
  - $\sigma(X, Y)$ is submodular in X when fixing Y, or vice versa
  - but the interaction of X and Y makes it harder --- need both X and Y to generate influence spread --- non-submodular behavior
- Hardness: NP-hard to approximate to any constant factor in general graphs
  - reduced from a k-prover system
- Algorithm: When M (weighted bi-adjacency matrix for provider-consumer bipartite graph) has constant rank, polynomial-time algorithm with constant approx. factor $\left(1 - \frac{1}{e} - \varepsilon\right)^3$
  - constant rank assumption:
    - M dictated by a small number of provider/consumer features
    - often used in recommender systems

The algorithm

- Main observation:

$\Pr[v \text{ activated}] = 1 - \prod_{u\in\Gamma(v)}(1 - M_{(u\to v)}x_u) \approx 1 - e^{-\sum_{u\in\Gamma(v)} M_{(u\to v)}x_u}$ (*)

$(1-1/e)$ factor

- Linear term in the exponent allows us to use constant rank
  - (polynomially sized) ε-net S_ε over all vectors $\mathbf{s} \approx M\mathbf{x}$
    - (fixing $\mathbf{s}$, we know from (*) the probability that each consumer is activated by the providers)

ε error

  - for each choice of $\mathbf{s}$

$(1-1/e-\varepsilon)$ factor

    - pick feasible $\mathbf{y_s}$ (indicator vector of V) that (approximately) maximizes the spread (standard submodular maximization)
    - Pick feasible $\mathbf{x}_{\mathbf{y_s}}$ (indicator vector of U) that (approximately) maximizes the spread, given fixed $\mathbf{y_s}$

$(1-1/e-\varepsilon)$ factor

  - output $(\mathbf{x}_{\mathbf{y_s}}, \mathbf{y_s})$ that maximizes $\sigma(\mathbf{x}_{\mathbf{y_s}}, \mathbf{y_s})$, among all $\mathbf{s} \in S_\varepsilon$

Contributions

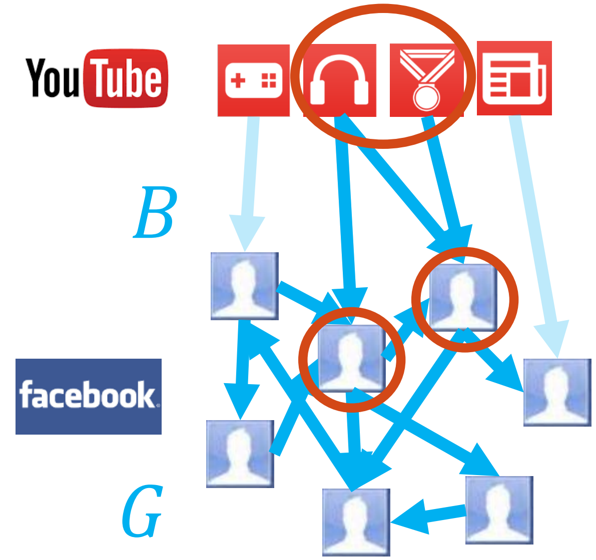

- Conceptual: proposes amphibious marketing
  - using data from both content providers and social networks
- Technical:
  - hardness reduction from k-prover systems
  - constant rank assumption to deal with non-submodularity

Open problems

- Is the approximation factor $\left(1 - \frac{1}{e} - \varepsilon\right)^3$ tight?
- Handling non-submodularity:
  - Can constant rank assumption be applied to other context?
  - Is there other reasonable assumptions on (weighted) influence networks?
- Better models for "amphibious marketing"?
  - learning? privacy? incentives?

Comparative Influence Diffusion: From Competition to Complementarity

with Wei Lu and Laks Lakshmanan

Competition and complementarity

- Many competitions in the market

vs

McDonald's vs BURGER KING

- Also many complementarity and cooperation

and Office


and


and DELL

Influence diffusion with competition and complementarity

- Most existing research focus on single-item diffusion or pure competitive diffusion
- Can we cover the diffusion of competitive and/or complementary items in a single unified model?

Comparative IC (Com-IC) model

- Consider two items A and B
- Comparative: user compares two items
  - covers both competition and complementarity, partial or complete
- Social graph $G = (V, E)$
  - edge probabilities $p(u, v)$: probability that u will pass information about A or B to v
    - open once for both items
- Node adoption states and their transitions
  - four states for each item:
    - idle, suspended, adopted, rejected

  - four global adoption probability (GAP) parameters:
    - $q_{A|\emptyset}, q_{A|B}, q_{B|\emptyset}, q_{B|A}$

Node level state transition

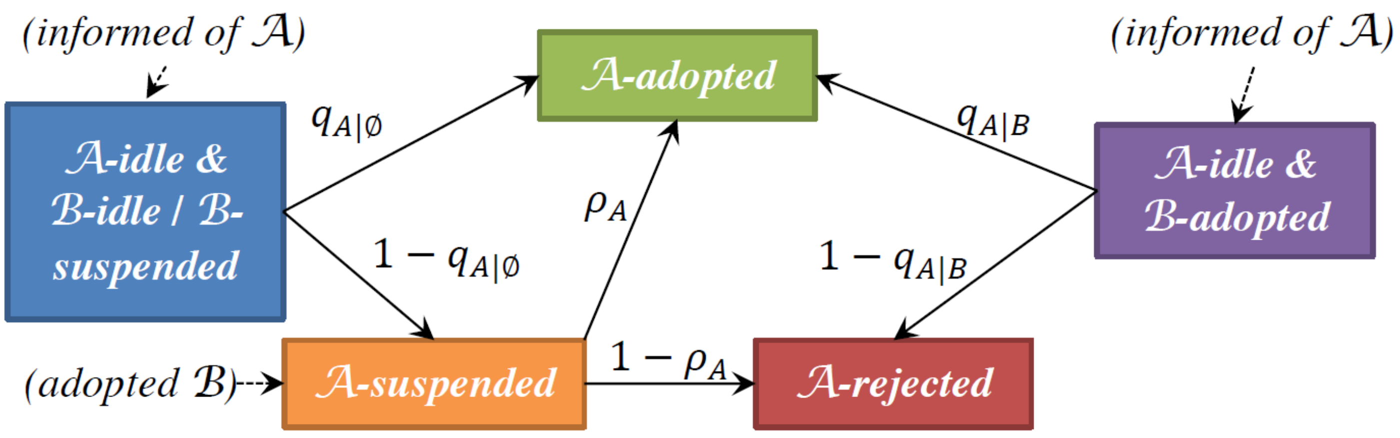

- Principle
  - If not adopt B and get informed about A, use $q_{A|\emptyset}$ to test A adoption
  - If already adopted B and get A-informed, use $q_{A|B}$
  - If failed adopting A initially (becoming A-suspended) and later adopt B, reconsider A with probability ρ_A
    - $\rho_A = \frac{\max\{q_{A|B}-q_{A|\emptyset},0\}}{1-q_{A|\emptyset}}$
    - only reconsider when $q_{A|B} \geq q_{A|\emptyset}$, and overall A adoption probability is $q_{A|B}$

The competition to complementarity spectrum

- Mutual competition:
  - $q_{A|B} \leq q_{A|\emptyset}$ and $q_{B|A} \leq q_{B|\emptyset}$
  - pure competition: $q_{A|B} = 0$ and $q_{B|A} = 0$
- Mutual complementarity:
  - $q_{A|B} \geq q_{A|\emptyset}$ and $q_{B|A} \geq q_{B|\emptyset}$
  - perfect complementarity: $q_{A|B} = 1$ and $q_{B|A} = 1$
- Mutual indifference:
  - $q_{A|B} = q_{A|\emptyset}$ and $q_{B|A} = q_{B|\emptyset}$
- One way complementarity:
  - $q_{A|B} \geq q_{A|\emptyset}$ and $q_{B|A} = q_{B|\emptyset}$

Submodularity on complementary cases

- influence spread $\sigma_A(S_A, S_B)$: expected number of A-adopted nodes with A seed set S_A and B seed set S_B
  - self-submodularity: fix S_B, σ_A changes on S_A
  - cross-submodularity: fix S_A, σ_A changes on S_B
- Submodularity only holds in some sub-cases
  - self-submodularity holds for one-way complementarity: $q_{A|B} \geq q_{A|\emptyset}$ and $q_{B|A} = q_{B|\emptyset}$
  - cross-submodularity holds when A perfectly complements B: $q_{B|A} = 1$

Influence maximization on complementary cases

- SelfInfMax: For a fixed set S_B, find S_A of size k to maximize $\sigma_A(S_A, S_B)$
- CompInfMax: For a fixed set S_A, find S_B of size k to maximize $\sigma_A(S_A, S_B) - \sigma_A(S_A, \emptyset)$
- Our results:
  - for submodular cases: design fast approximation algorithm based on Reverse-Reachable sets
  - for non-submodular cases: finder upper/lower bounded submodular cases, and use sandwich approximation

Experimental evaluation

$\mathcal{A}$	$\mathcal{B}$	$q_{\mathcal{A}\vert\emptyset}$	$q_{\mathcal{A}\vert\mathcal{B}}$	$q_{\mathcal{B}\vert\emptyset}$	$q_{\mathcal{B}\vert\mathcal{A}}$
Monster Inc.	*Shrek*	.88	.92	.92	.96
Gone in 60 Seconds	*Armageddon*	.63	.77	.67	.82
Prisoner of Azkaban	*What a Girl Wants*	.85	.84	.66	.67
Shrek	*Fast and Furious*	.92	.94	.80	.79

Table 2: Selected GAPs learned for movies from *Flixster*

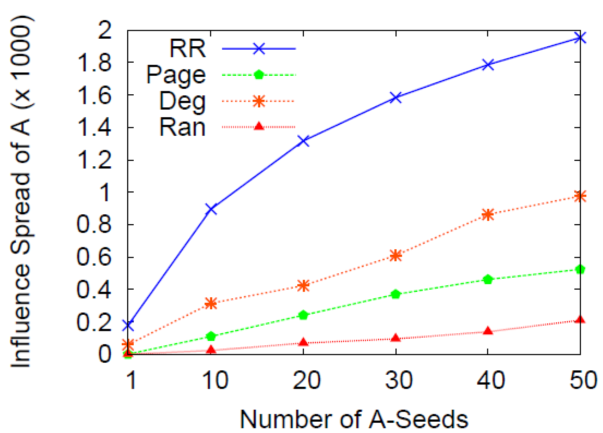

SelfInfMax on Flixster

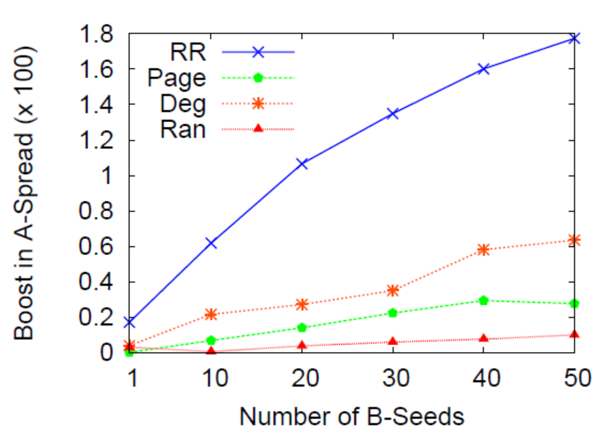

CompInfMax on Flixster

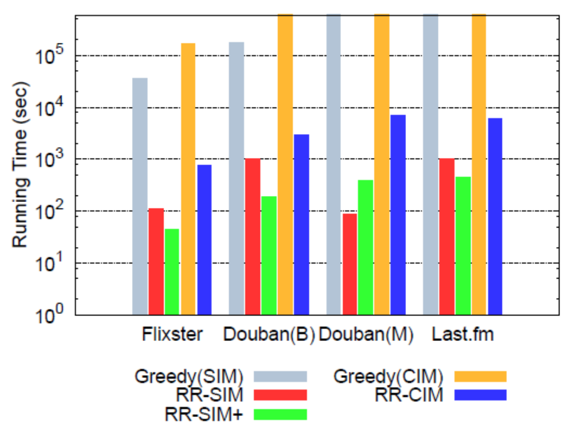

Running time

Contribution

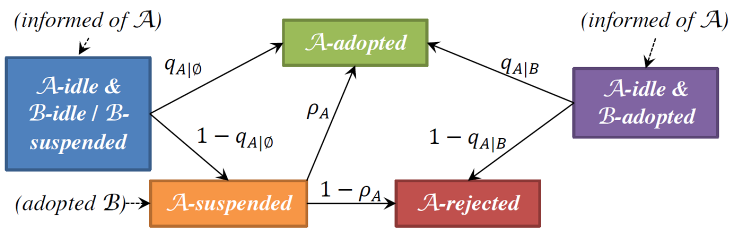

- Conceptual: propose a unified diffusion model covering both competition and complementarity
- Technical:
  - new problems arises from the new model
  - self- and cross-submodularity analysis
  - generalize RR-set approach and design fast approximation algorithms for SelfInfMax and CompInfMax
  - sandwich approximation dealing with non-submodularity

Open problems

- Can SelfInfMax and CompInfMax be made near-linear time?
- Can we fully characterize monotonicity and submodularity for the entire GAP space?
- Can we efficiently generalize Com-IC to multiple items?
- How to efficiently learn GAP parameters?
- What are other interesting problems in the Com-IC model?

Conclusion

CSI still in its early stage

- Many models and problems still need to be studied
  - non-binary, non-progressive models
  - dealing with dynamic graphs
- Influence analysis and learning is still a big challenge
  - data is still not big enough!
    - too sparse, too noisy, non-critical
  - need smart methods in learning influence model parameters
  - need robust optimization methods dealing with uncertainty in the model
  - combine online learning with influence maximization

Grand challenge

- Understand from data the true peer influence and viral diffusion scenarios, online and offline
- Apply social influence research to explain, predict, and control influence and viral phenomena
- Network and diffusion dynamics would be focus of network science in the next decade

Further resources

Search "Wei Chen Microsoft"

- Monograph: "Information and Influence Propagation in Social Networks", Morgan & Claypool, 2013
- KDD'12 tutorial on influence spread in social networks
- my papers and talk slides

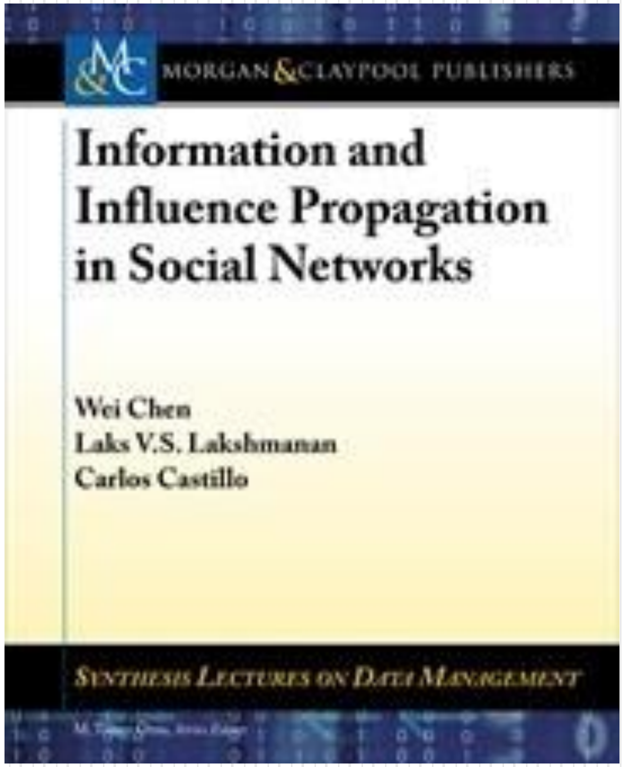

Acknowledgments to my CSI collaborators

Yajun Wang, Siyu Yang, Chi Wang, Yifei Yuan, Li Zhang, Alex Collins, Rachel Cummings, Te Ke, Zhenming Liu, David Rincon, Xiaorui Sun, Wei Wei, Tao Sun, Ming Zhang, Chin-Yew Lin, Pinyan Lu, Bo Tang, Zeyuan Allen-Zhu, Xinran He, Guojie Song, Qingye Jiang, Wei Lu, Ning Zhang, Kyomin Jung, Wooran Heo, Yanhua Li, Zhi-Li Zhang, Yang Yuan, De-Nian Yang, Hui-Ju Hung, Wang-Chien Lee, Sejeong Kwon, Meeyoung Cha, Laks Lakshmanan, Carlos Castillo, Peng Zhang, Xiaoming Sun, Jialin Zhang, Jing Zhang, Zhanpeng Fang, Jie Tang, Fu Li, Tian Lin, Aviad Rubinstein, Cheng Yang

Thank you!