Influence Maximization with Bandits

Sharan Vaswani, Laks.V.S.Lakshmanan, Mark Schmidt
Department of Computer Science, University of British Columbia

Motivation

- **Viral marketing** uses a social network to spread awareness about a product.

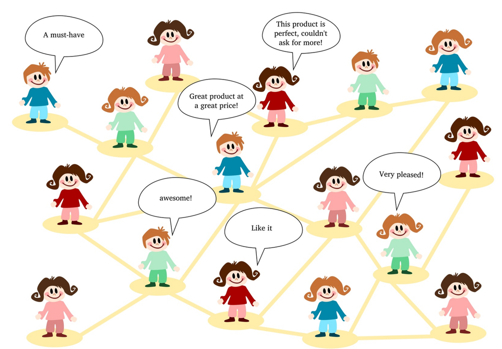

- **Influence Maximization (IM):**
  - Select a fixed number of 'influential' users (seeds) to give free products or discounts.
  - Try to maximize the number of people who become aware of the product (spread),

$$S^* = \text{argmax}_{|S| \leq k} \sigma_D(S).$$

where S are the seeds, k is the budget, and $\sigma_D(S)$ is expected spread under stochastic diffusion model D.

- **Limitations of existing methods**:
  - assume you know the pairwise influence probabilities (could be hard to obtain in practice).
  - assume edge-level feedback: know which user influenced each other user (often not realistic).
- **Our contributions**:
  - Formulate as combinatorial multiarmed bandit problem.
  - Aim to minimize regret as a new marketer learns the influence probabilities.
    - Leads to classic exploration vs. exploitation trade-off.
  - Consider node-level feedback: you only need to know who was influenced.

Background on Independent Cascade and Multiarmed Bandits

- **Independent Cascade (IC) Model:**
  - Starting from seeds, influenced nodes get one chance to influence their neigbhours.
  - Succeed with probability $p_{u,v}$ (*live* edge) and otherwise fail (*dead*).
  - Newly-influence nodes can influence their neighbours.
- **Multiarmed and Combinatorial Multiarmed Bandits:**
  - Each of m arms has reward distribution with unknown mean μ.
  - Standard framework: in round t you choose one arm i and obtain reward $r_{i,t}$.
  - Combinatorial framework: in round t you choose a subset of arms A and reward is function of these arms.

Mapping Influence Maximization to Combinatorial Multiarmed Bandits

We can write influence maximization in combinatorial multiarmed bandit framework:

CMAB	Symbol	Mapping to IM
Base arm	i	Edge (u, v)
Superarm	E_S	Union of outgoing edges from nodes in set S
Reward for arm i in round s	$X_{i,s}$	Status (live/dead) for edge (u, v)
Mean of distribution for arm i	μ_i	Influence probability $p_{u,v}$
No. of times i is triggered in s rounds	$T_{i,s}$	#times u becomes active in s diffusions
Reward in round s	r_s	Spread $\bar{\sigma}$ in the s^{th} IM attempt

Algorithm 1: CMAB FRAMEWORK FOR IM (Graph $G = (V, E)$, budget k, Feedback mechanism M, Algorithm $\mathcal{A}$)

```
Initialize μ⃗ ;
∀i initialize T_i = 0 ;
for s = 1 → T do
    IS-EXPLOIT is a boolean set by algorithm A ;
    if IS-EXPLOIT then
        E_S = EXPLOIT(G,μ⃗,O,k)
    else
        E_S = EXPLORE(G,k)
    Play the superarm E_S and observe the diffusion cascade c ;
    μ⃗ = UPDATE(c,M) ;
```

Edge-Level and Node-Level Feedback Mechansim

- **Edge Level Feedback (EL):**
  - Assumes you can view the status of edge i.
  - Simple update of influence probabilities: $\hat{\mu}_i = \frac{\sum_{s=1}^{t} X_{i,s}}{T_{i,t}}$.
  - Often not realistic: we can see whether user adopted a product, not who did/didn't influence them.
- **Node Level Feedback (NL):**
  - Assumes you can view the status of each node.
  - More realistic: typically easy to observe in network.
  - But updating influence probabilities requires assigning credit.

Bounding the error for node-level credit assignment

- We consider a simple heuristic credit assignment mechanism for node-level feedback:
  - Each active node v randomly chooses one of its active parents u, and assigns full credit to edge (u, v).
  - Makes node-level feedback effective in typical social networks where influence probabilities are typically low.

Theorem

Let p_{min} and p_{max} be the minimum and maximum true influence probabilities in the network. Consider a particular cascade c and any active node v with K_c active parents. The failure probability ρ under our node-level feedback credit assignment scheme for node v satsifies

$$\rho \leq \frac{1}{K_c}(1 - p_{min})\left(1 - \prod_{k=1,k\neq i}^{K_c}[1 - p_{max}]\right) + \left(1 - \frac{1}{K_c}\right)p_{max}. \quad (1)$$

Suppose $\hat{\mu}_i^E$ and $\hat{\mu}_i^N$ are the inferred influence probabilities for the edge corresponding to arm i using edge-level and node-level feedback respectively. Then the relative error in the learned influence probability is given by:

$$\left|\frac{\hat{\mu}_i^N - \hat{\mu}_i^E}{\hat{\mu}_i^E}\right| = \rho\left|\frac{1}{\hat{\mu}_i^E} - 2\right| \quad (2)$$

Regret Minimization Algorithms

- **Upper Confidence Bound (UCB):** combinatorial UCB maintains an overestimate $\overline{\mu}_i$ of the mean estimates $\hat{\mu}_i$.
- **Pure Exploitation:** performs exploitation in every round.
- **ϵ-Greedy:** exploration with probability ϵ_s and exploitation with probability $1 - \epsilon_s$.

Numerical Experiments

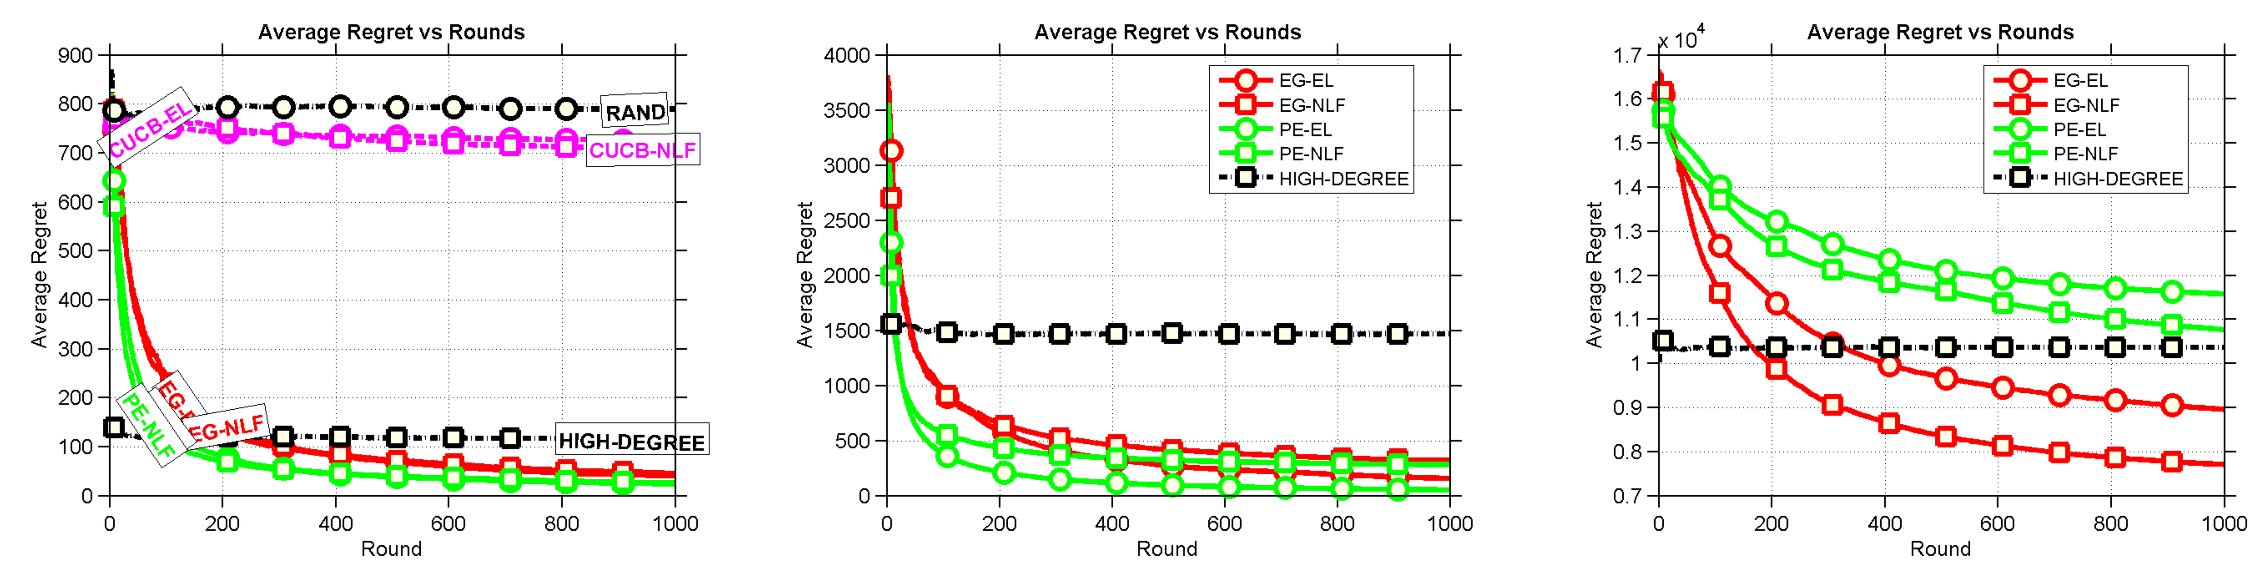

(a) NetHEPT (b) Flixster (c) Flickr

- Pure Exploitation (PE), ϵ-Greedy (EG) are effective and able to decrease the regret across all datasets.
- Node Level feedback (NL) has results comparable to Edge Level feedback (EL) for all algorithms across datasets.

Related Work

- Regret analysis under UCB for IM (Chen et al., 2014).
- Multiple IM attempts to maximize the number of distinct active nodes across rounds (Lei et al., 2015).