Model-Independent Online Learning for Influence Maximization

Sharan Vaswani [1] Branislav Kveton [2] Zheng Wen [2] Mohammad Ghavamzadeh [3] Laks.V.S.Lakshmanan [1] Mark Schmidt [1]
[1] University of British Columbia [2]Adobe Research [3]Deepmind

Influence Maximization

- **Underlying principle:** Influence propagates through word of mouth in a social network.
- **Idea:** Give discounts to influential users who will trigger off word-of-mouth epidemics.
- **Aim:** Find a subset of users ('seed set') who will influence maximum people to become aware of a product:
- **Challenge 1:** IM is not robust to the choice of the diffusion model (Du, 2014) nor its model parameters (Goyal, 2010).
- **Challenge 2:** It is difficult for a new marketer to have data to learn the large number of model parameters.
- **Summary of paper:**
  - Develop a model-independent parametrization for IM and a corresponding surrogate objective function.
  - Propose and analyze a UCB based algorithm for model-independent online IM.
  - Propose a scalable linear parametrization and empirically verify its effectiveness.

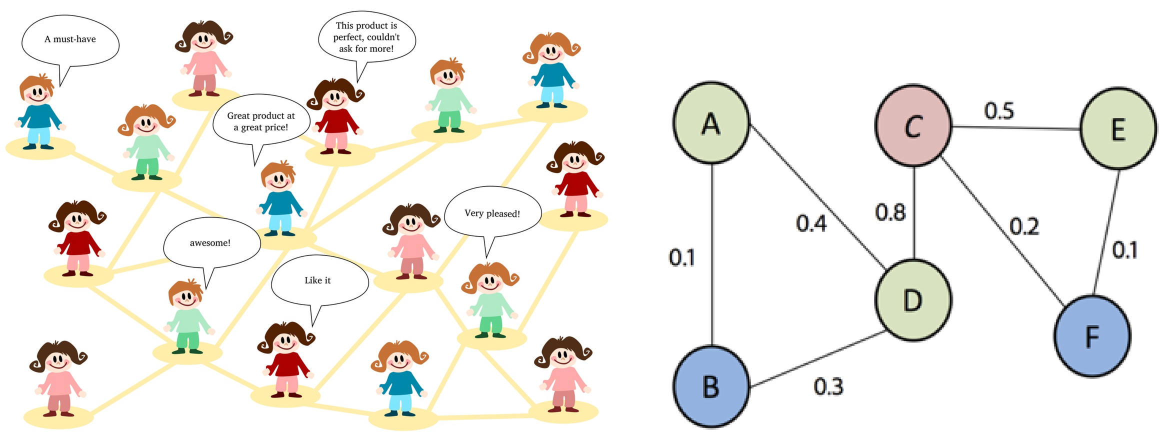

(a) Information diffusion in a social network (b) Input for an IM problem

- **Input:** Graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$; Set of feasible seed sets $\mathcal{C} \subseteq \{\mathcal{S} \subseteq \mathcal{V} : |\mathcal{S}| \leq K\}$; Stochastic diffusion model $\mathcal{D}$.
- **Formal objective:** Find $\mathcal{S}^* \in \arg\max_{\mathcal{S} \in \mathcal{C}} F(\mathcal{S})$ where $F(\mathcal{S}) = \sum_{v \in \mathcal{V}} F(\mathcal{S}, v)$.

Model-Independent Formulation

- **Assumption 1:** The diffusion model is progressive i.e. $F(\mathcal{S})$ is monotonic in $\mathcal{S}$.
- **Key idea:** Parametrize the problem in terms of pairwise reachability probabilities $p^*_{u,v} = F(\{u\}, v)$.
- **Surrogate objective:** Find $\widetilde{\mathcal{S}}$ s.t. $\widetilde{\mathcal{S}} \in \arg\max_{\mathcal{S} \in \mathcal{C}} \left[\sum_{v \in \mathcal{V}} \left(\max_{u \in \mathcal{S}} p^*_{u,v}\right)\right]$.
- **Advantage:** Surrogate objective is submodular irrespective of the diffusion model.

Theorem

For any graph $\mathcal{G}$, seed set $\mathcal{S} \in \mathcal{C}$, and diffusion model $\mathcal{D}$ satisfying Assumption 1,

1. $f(\mathcal{S}, p^*) \leq F(\mathcal{S})$,
2. *If $F(\mathcal{S})$ is submodular in $\mathcal{S}$ and $\rho = f(\widetilde{\mathcal{S}}, p^*)/F(\mathcal{S}^*)$, then $1/K \leq \rho \leq 1$.*

Online Influence Maximization

- **Setting:** New marketer who has no past data to learn the reachability probabilities.
- **Idea:** Perform IM while simultaneously learning parameters through trial and error across multiple rounds.
- **Protocol:** For $t = 1$ to T:
  1. Select seed set: $\mathcal{S}_t \leftarrow \texttt{ORACLE}(\mathcal{G}, \mathcal{C}, \overline{p})$.
  2. Diffusion occurs according to an underlying diffusion model.
  3. Observe semi-bandit feedback: $\forall u \in \mathcal{S}_t$, get pairwise influence feedback $\mathbf{y}_{u,t}$.
  4. Update parameter estimates $\overline{p}_{u,v}$.
- **Challenges:**
  1. Learn n^2 reachability probabilities.
  2. Choose $\mathcal{S}_t$ to trade off exploration and exploitation.
- **Linear parametrization:** For all $u, v \in \mathcal{V}$, $p^*_{u,v} \approx \langle \boldsymbol{\theta}^*_u, \mathbf{x}_v \rangle$.
  1. Reduces the number of parameters to $O(dn)$.
  2. In each round, mean estimates of $\overline{p}_{u,v}$ can be updated efficiently by solving K regression problems.

DILinUCB Algorithm

- Upper confidence bound (UCB) based algorithm: If $\widehat{\boldsymbol{\theta}}_{u,t}$ and $\Sigma_{u,t}$ is the mean estimate and covariance matrix for the regression problem for node u at round t, then $\overline{p}_{u,v} = \text{Proj}_{[0,1]}\left[\langle \widehat{\boldsymbol{\theta}}_{u,t} \mathbf{x}_v \rangle + c\|\mathbf{x}_v\|_{\Sigma^{-1}_{u,t}}\right]$.
- **Computational Complexity:** $O(Knd^2)$ for updating UCBs and $O(Kn)$ for oracle computation.

Regret Bound

- **Performance Metric:** Scaled regret: $R^{\kappa}(T) = T \cdot F(\mathcal{S}^*) - \frac{1}{\kappa}\mathbb{E}\left[\sum_{t=1}^{T} F(\mathcal{S}_t)\right]$.
- With linear generalization: $R^{\rho\alpha}(T) = \widetilde{O}(n^2 d\sqrt{KT}/(\alpha\rho))$.
- Without linear generalization: $R^{\rho\alpha}(T) = \widetilde{O}(n^{2.5}\sqrt{KT}/(\alpha\rho))$.
- Model-independent regret bound with near optimal dependence on T, d and K.

Experiments

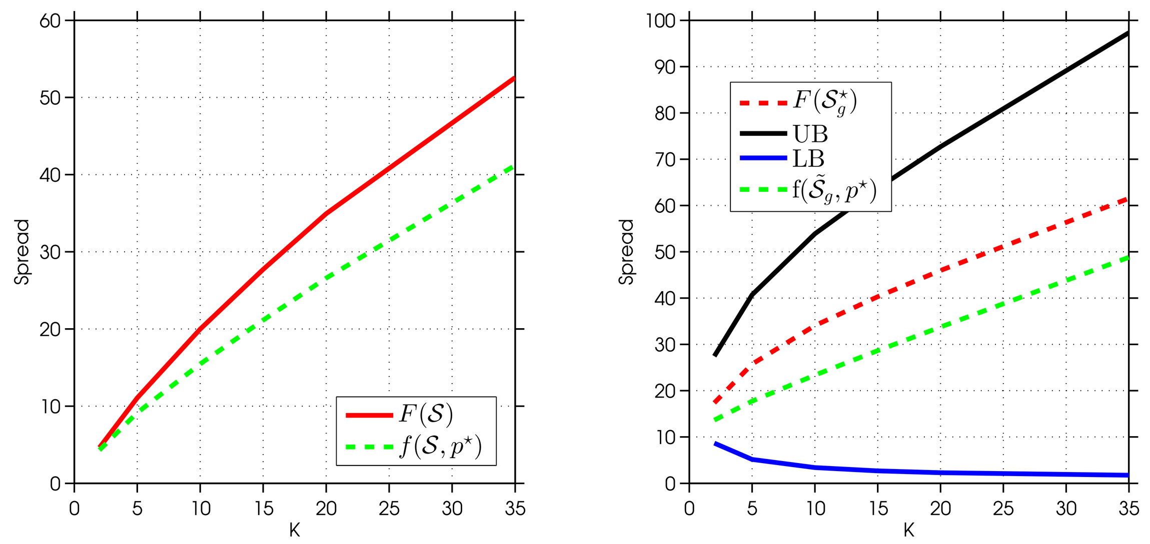

(c) (d)

Figure: Experimental verification of surrogate objective.

Algorithms compared: (1) CUCB (Chen, 2016) [CUCB(K)] (2) DILinUCB without linear generalization [$TAB(K)$] (3) DILinUCB with linear generalization [$I(K, d)$] (4) DILinUCB with linear generalization + Laplacian regularization [$L(K, d)$].

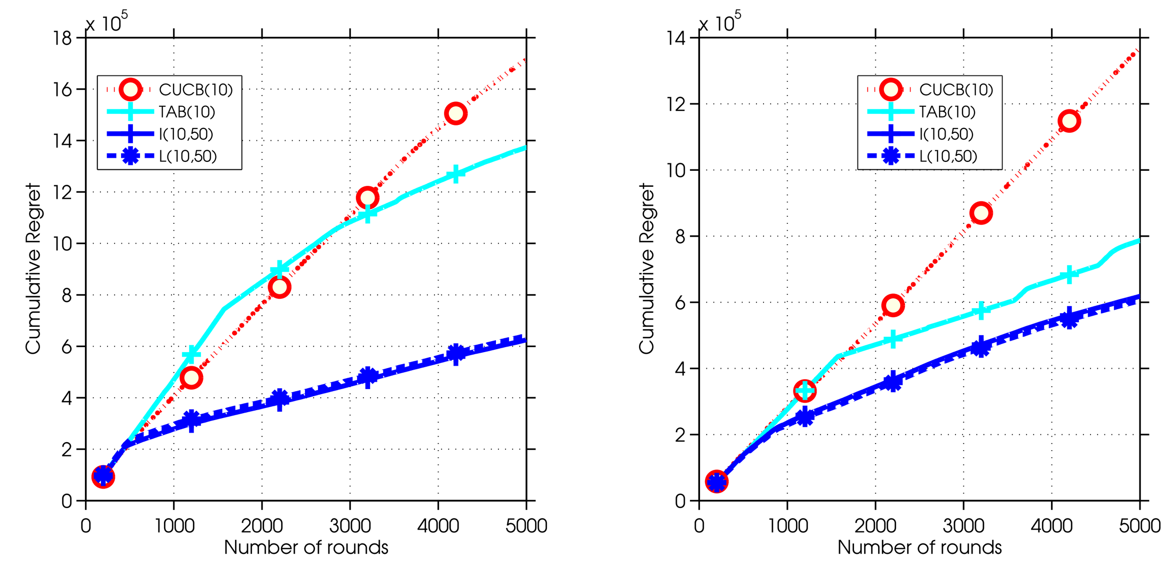

(a) IC Model (b) LT Model

Figure: Comparing DILinUCB and CUCB on the Facebook subgraph with $K = 10$.

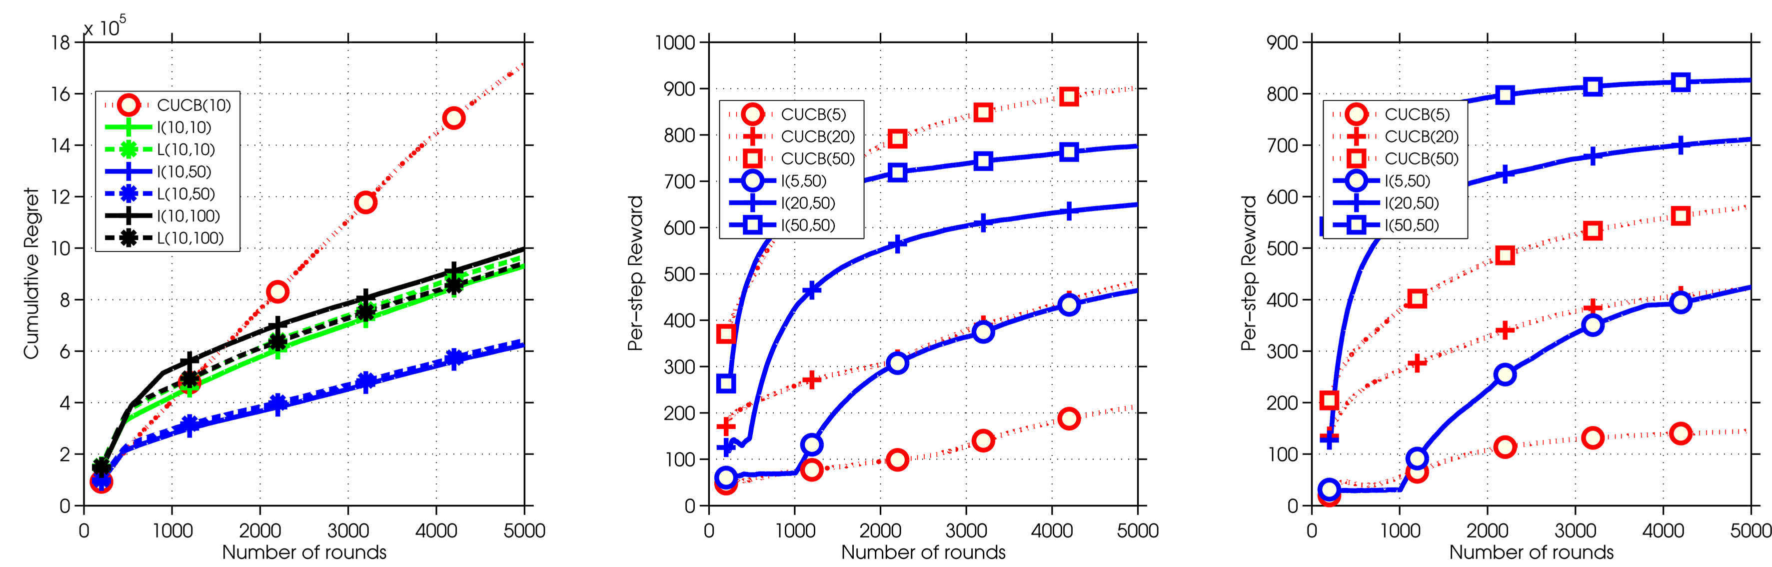

(a) Effect of d in IC (b) Effect of K in IC (c) Effect of K in LT

Figure: Effects of varying d or K.

Future Work

- Extend the framework to different feedback models and bandit algorithms.
- Generalization across source nodes for better statistical efficiency and improved regret bounds.