Social Data Management Probabilistic Graphs and Influence Algorithms

Silviu Maniu[1]

December 10th, 2018

[1]Université Paris-Sud

Uncertain Graphs

Graphs: a natural way to represent data in various domains

- transport data: road, air links between locations
- social networks: relationships between humans, citation networks
- interactions between proteins: contacts due to biochemical processes

Uncertain Graphs

Graphs: a natural way to represent data in various domains

- transport data: road, air links between locations
- social networks: relationships between humans, citation networks
- interactions between proteins: contacts due to biochemical processes

For all the above examples, the links are not exact. (*Why?*)

(Deterministic) Graphs

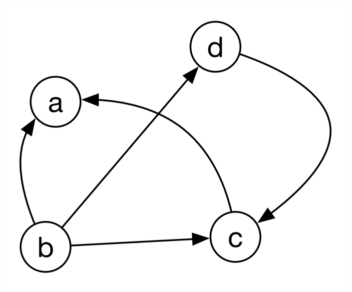

A graph $G = (V, E)$ is formed of

- a set V of vertices (nodes)
- a set $E \subseteq V \times V$, of edges

Uncertain Graphs

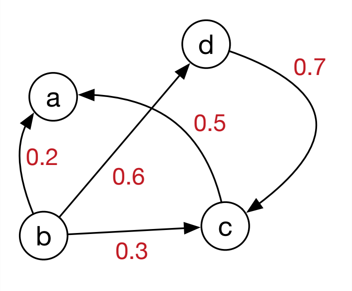

An uncertain graph $\mathcal{G} = (V, E, p)$ is formed of

- a set V of vertices (nodes)
- a set $E \subseteq V \times V$, of edges
- a function $p : E \to [0, 1]$, representing the probability p_e that the edge $e \in E$ exists or not

What are the possible worlds and their probability for this model?

Uncertain Graphs: Possible Worlds

A possible world of $\mathcal{G}$, denoted $G \sqsubseteq \mathcal{G}$ is a *deterministic* graph $G = (V, E_G)$ where each $e \in E_G$ is chosen from E

Uncertain Graphs: Possible Worlds

A possible world of $\mathcal{G}$, denoted $G \sqsubseteq \mathcal{G}$ is a *deterministic* graph $G = (V, E_G)$ where each $e \in E_G$ is chosen from E

The probability of G is:

$$\Pr[G] = \prod_{e \in E_G} p_e \prod_{e \in E \setminus E_G} (1 - p_e)$$

How many possible worlds are there?

Uncertain Graphs: Other models

Other models are possible:

- each edge is replaced by a distribution of weights – instead of choosing if the edge exists or not, a possible world is an instantiation of weights
- each edge has a formula of events, capturing correlations
- probabilities can be on nodes also – equivalent to the edge model (*Why?*)

Queries on Uncertain Graphs

Generally, the queries we want to answer are distance queries:

- the reachability or reliability query – get the probability that two nodes s and t are connected

Queries on Uncertain Graphs

Generally, the queries we want to answer are distance queries:

- the reachability or reliability query – get the probability that two nodes s and t are connected
- queries on the distance distribution:

$$p_{s,t}(d) = \sum_{G | d_G(s,t) = d} \Pr[G]$$

Queries on Uncertain Graphs

Generally, the queries we want to answer are distance queries:

- the reachability or reliability query – get the probability that two nodes s and t are connected
- queries on the distance distribution:

$$p_{s,t}(d) = \sum_{G | d_G(s,t)=d} \Pr[G]$$

Multiple uses of distance queries:

- link prediction, social search, travel estimation

Queries on Uncertain Graphs

Generally, the queries we want to answer are distance queries:

- the reachability or reliability query – get the probability that two nodes s and t are connected

Queries on Uncertain Graphs

Generally, the queries we want to answer are distance queries:

- the reachability or reliability query – get the probability that two nodes s and t are connected
- queries on the distance distribution:

$$p_{s,t}(d) = \sum_{G | d_G(s,t) = d} \Pr[G]$$

Queries on Uncertain Graphs

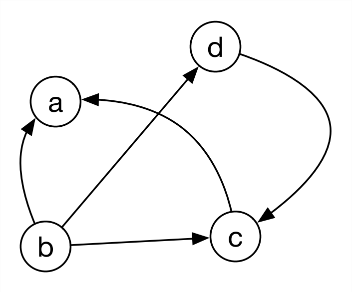

What is the distance (in hops) between b and a ?

Queries on Uncertain Graphs

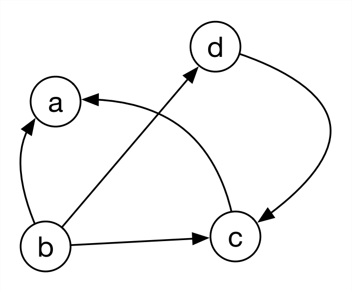

What is the distance (in hops) between b and a ?

- BFS search (or Dijkstra's algorithms) finds the edge $b \rightarrow a$
- the cost is $O(E)$ (linear in the size of the graph)

Queries on Uncertain Graphs

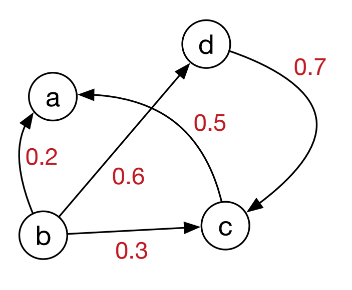

What is the distance (in hops) between *b* and *a* ?

Queries on Uncertain Graphs

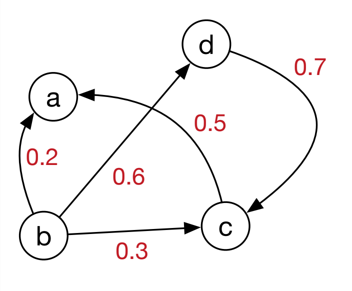

What is the distance (in hops) between b and a ?

- the edge $b \to a$ does not appear in all possible worlds:

$$p_{b,a}(1) = p(b \to a)$$

Queries on Uncertain Graphs

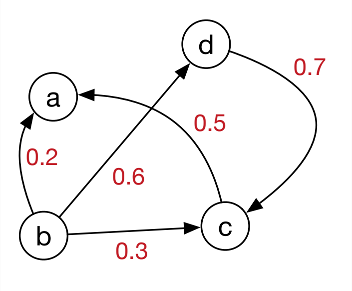

What is the distance (in hops) between b and a ?

- the edge $b \rightarrow a$ does not appear in all possible worlds:

$$p_{b,a}(1) = p(b \rightarrow a)$$

- there are two possible paths of distance 2 ($b \rightarrow c \rightarrow a$) and 3 ($b \rightarrow d \rightarrow c \rightarrow a$)

$$p_{b,a}(1) = (1 - p_{b,a}(1)) \times p(b \rightarrow c \rightarrow a)$$

Queries on Uncertain Graphs

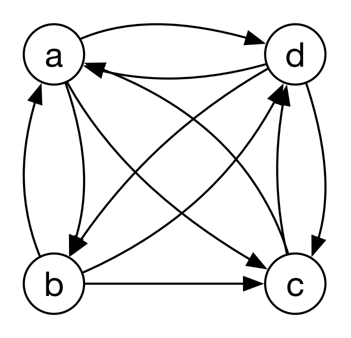

What is the distance (in hops) between b and a ?

Queries on Uncertain Graphs

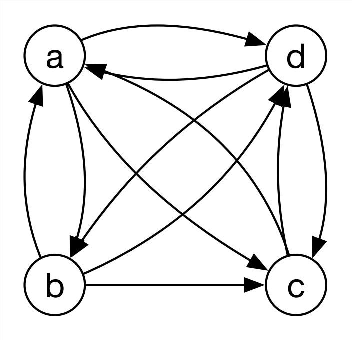

What is the distance (in hops) between *b* and *a* ?

- the number of paths is exponential in the size of the graph
- specifically, there are 3! paths

Queries on Uncertain Graphs

Distance query answering in uncertain graphs is at least as hard as in relational databases (*logical formulas* of paths; the number of which can be exponential)

Queries on Uncertain Graphs

Distance query answering in uncertain graphs is at least as hard as in relational databases (*logical formulas* of paths; the number of which can be exponential)

Computing the reachability probability (i.e, computing the probability of there being a path between a source and a target) is known to be #P hard [Valiant, SIAM J. Comp, 1979]

Computing Answers to Distance Queries on Probabilistic Graphs

Distance estimations in uncertain graphs can be approximated via Monte Carlo sampling

Computing Answers to Distance Queries on Probabilistic Graphs

Distance estimations in uncertain graphs can be approximated via Monte Carlo sampling

1. generate sampled graphs for r rounds (is this the optimal way for an s, t distance estimation?)
2. compute the desired measure (reachability probability, distance distributions) by averaging results

Computing Answers to Distance Queries on Probabilistic Graphs

Distance estimations in uncertain graphs can be approximated via Monte Carlo sampling

1. generate sampled graphs for r rounds (is this the optimal way for an s, t distance estimation?)
2. compute the desired measure (reachability probability, distance distributions) by averaging results

Same issue: *how many rounds?*

Number of Samples: Median Distance

Median distance:

$$d_M(s,t) = \arg\max_D \left\{ \sum_{d=0}^{D} p_{s,t}(d) \leqslant \frac{1}{2} \right\}$$

Number of Samples: Median Distance

Median distance:

$$d_M(s,t) = \arg\max_D \left\{ \sum_{d=0}^{D} p_{s,t}(d) \leqslant \frac{1}{2} \right\}$$

Let μ be the real median, and α and β values $\pm\epsilon N$ away from μ. Then for:

$$r > \frac{c}{\epsilon^2} \log(\frac{2}{\delta})$$

and a good choice of c:

$$\Pr(\hat{\mu} \in [\alpha, \beta]) > 1 - \delta$$

Number of Samples: Expected Distance

Expected reliable distance (generalization of reliability):

$$d_{\mathrm{ER}}(s, t) = \sum_{d|d<\infty} d \cdot \frac{p_{s,t}(d)}{1 - p_{s,t}(\infty)}$$

Number of Samples: Expected Distance

Expected reliable distance (generalization of reliability):

$$d_{\mathsf{ER}}(s,t) = \sum_{d \mid d < \infty} d \cdot \frac{p_{s,t}(d)}{1 - p_{s,t}(\infty)}$$

By estimating the connectivity ρ, we need to sample at least:

$$r \geqslant \max\left\{\frac{3}{\epsilon^2\rho}, \frac{(n-1)^2}{2\epsilon^2}\right\} \cdot \log\left(\frac{2}{\delta}\right)$$

for an (ϵ, δ) approximation.

Number of Samples In Reality

The number of needed samples can be surprisingly low (but it depends on the actual probabilities)

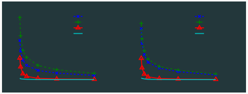

Sampling Graphs

Generating the entirety of the graph G_i for each round $i < r$ is not optimal

Sampling Graphs

Generating the entirety of the graph G_i for each round $i < r$ is not optimal

- we do not need to estimate the entire graph G_i
- we can start from s and do a BFS or Dijkstra search by sampling only the outgoing edges
- based on the generated outgoing edges, we re-do the computation for each generated outgoing node, until we find t

Example: Median Distance k-NN

k-NN (k nearest neighbours) – finding the k nodes from s the "closest" by some measure

- let us consider the median distance (reminder: it is the highest distance in the distribution that has mass less or equal to 0.5)

Example: Median Distance k-NN

k-NN (k nearest neighbours) – finding the k nodes from s the "closest" by some measure

- let us consider the median distance (reminder: it is the highest distance in the distribution that has mass less or equal to 0.5)

We only care about the top-k nodes, and not their values, and we do not want to evaluate all the graph if possible

- we can evaluate a truncated distribution up to a distance D

$$p_{D,s,t}(d) = \begin{cases} p_{s,t}(d) & \text{if} \quad d < D \\ \sum_{x=D}^{\infty} p_{s,t}(x) & \text{if} \quad d = D \\ 0 & \text{if} \quad d > D \end{cases}$$

- for any two nodes t_1, t_2, $d_{D,M}(s, t_1) < d_{D,M}(s, t_2)$ implies $d_M(s, t_1) < d_M(d, t_2)$

Example: Median Distance k-NN

Input: Probabilistic graph $\mathcal{G} = (V, E, P, W)$, node $s \in V$, number of samples r, number k, distance increment γ
Ouput: T_k, a result set of k nodes for the k-NN query

$T_k \leftarrow \emptyset$; $D \leftarrow 0$
Initiate r executions of Dijkstra from s
while $|T_k| < k$ **do**
  $D \leftarrow D + \gamma$
  for $i \leftarrow 1 : r$ **do**
    Continue visiting nodes in the i-th execution of Dijkstra until reaching distance D
    For each node $t \in V$ visited update the distribution $\tilde{\mathbf{p}}_{D,s,t}$ {Create the distribution $\tilde{\mathbf{p}}_{D,s,t}$ if t has never been visited before}
  end for
  for all nodes $t \notin T_k$ for which $\tilde{\mathbf{p}}_{D,s,t}$ exists **do**
    if median$(\tilde{\mathbf{p}}_{D,s,t}) < D$ **then**
      $T_k \leftarrow T_k \cup \{t\}$
    end if
  end for
end while

- start from a small distance D
- decide whether there are nodes to add to the k-NN set
- increase the distance, and "re-start" each sampled graph from the new distance

The algorithm does not need to visit all nodes

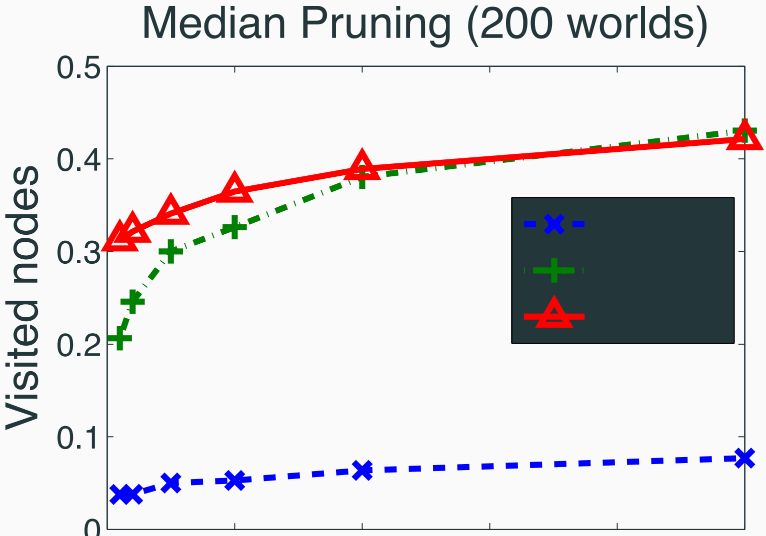

Table of contents

Social Influence

Social Influence: important problem in social network, with applications in marketing, computational advertising

Social Influence

Social Influence: important problem in social network, with applications in marketing, computational advertising

Objective: given a promotion budget of k social network users, maximize the expected influence spread given some influence or propagation model

Social Influence

Data Model: an uncertain graph $G(V, E, p)$

Data Model: an uncertain graph $G(V, E, p)$

- V and E are the social network
- p is, on each edge, the influence probability

Influence Spread via Cascades

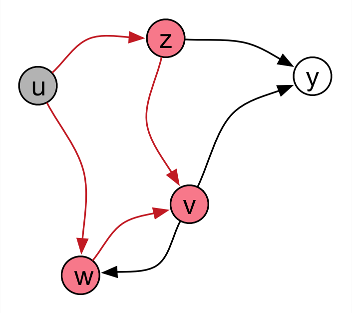

Independent Cascade Model:
discrete time model of propagation

1. at time 0, activate seed u
2. for a node i activated at time t: activate at time $t+1$ each neighbour v with probability p_{iv}
3. once a node is activated, it cannot be activated again or de-activated

Influence Spread via Cascades

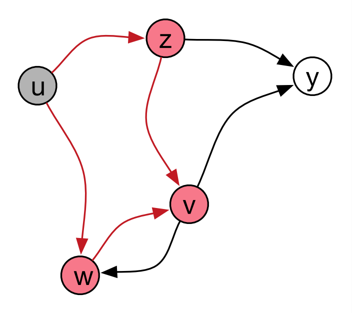

We wish to compute the expected spread from a seed seed set S, $\sigma(S)$

Influence Spread via Cascades

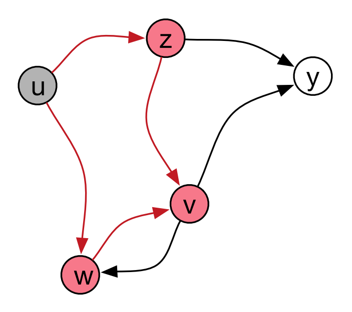

We wish to compute the expected spread from a seed seed set S, $\sigma(S)$

By linearity of expectation:

$$\sigma(u) = \sum_{v \in V} \Pr(u \to v)$$

- for a seed set S, more complicated
- same hardness as reachability

Maximizing the Influence

Influence maximization is computationally hard

Two sources of hardness:

1. computing $\sigma(S)$ is #P-hard (as we seen before, it is equivalent to reachability) – Monte Carlo with additive approximations
2. computing the selection of k seeds in S is NP-hard – maximization of a submodular function

Maximizing the Influence

Influence maximization is computationally hard

Two sources of hardness:

1. computing $\sigma(S)$ is #P-hard (as we seen before, it is equivalent to reachability) – Monte Carlo with additive approximations
2. computing the selection of k seeds in S is NP-hard – maximization of a submodular function

Submodular function: the influence spread is submodular:

$$\sigma(S \cup \{u\}) - \sigma(S) \geqslant \sigma(T \cup \{u\}) - \sigma(T) \quad \text{if} \quad S \subseteq T$$

Influence Maximization: Greedy Algorithm

We can obtain a $(1 - \frac{1}{\epsilon})$-approximation factor for influence maximization by using the greedy algorithm

Influence Maximization: Greedy Algorithm

We can obtain a $(1 - \frac{1}{\epsilon})$-approximation factor for influence maximization by using the greedy algorithm

Steps:

1. initialize $S = \emptyset$
2. choose the user u that maximizes $\sigma(S \cup \{u\}) - \sigma(S)$
3. $S = S \cup u$
4. repeat steps 2 and 3 k times
5. **return** S

Learning Propagation Probabilities

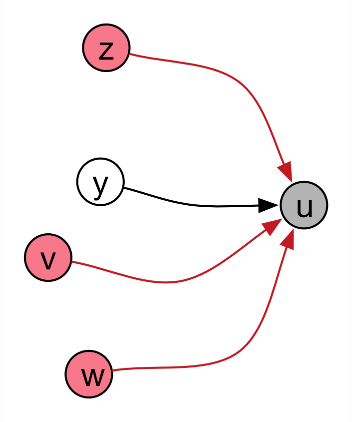

The probability that v is influenced by its neighbours

$$\Pr(v) = 1 - \prod_{u} (1 - p_{uv})$$

Learning Propagation Probabilities

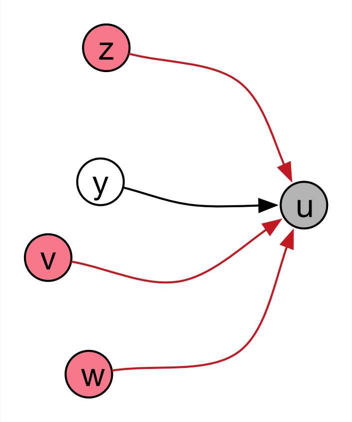

The probability that v is influenced by its neighbours

$$\Pr(v) = 1 - \prod_u (1 - p_{uv})$$

Given a log of actions $A = \{(\text{act}, u, v), \dots\}$:

1. maximum likelihood: $p_{vu} = \frac{A_{vu}}{A_v}$
2. Jaccard similarity: $p_{vu} = \frac{A_{vu}}{A_{u|v}}$

Acknowledgments

Figures in slides 16 and 20 are taken from [Potamias et al., 2010].

Ball, M. O. (1986).
Computational complexity of network reliability analysis: An overview.
IEEE Transactions on Reliability, 35(3).

Goyal, A., Bonchi, F., and Lakshmanan, L. V. (2010).
Learning influence probabilities in social networks.
In *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, WSDM '10.

References ii

Kempe, D., Kleinberg, J., and Tardos, E. (2003).
Maximizing the spread of influence through a social network.
In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '03.

Potamias, M., Bonchi, F., Gionis, A., and Kollios, G. (2010).
K-nearest neighbors in uncertain graphs.
Proc. VLDB Endow., 3(1-2).