~~Three~~ One result~~s~~ on recommendation systems

Prabhakar Raghavan, Google
Joint work with M. Bressan, S. Leucci, A. Panconesi and E. Terolli.

Recommendation systems (RS)

We have n users and m items.

Users buy items in a discrete-time process.

The item bought by a user at time t depends on purchases of other users before t, due to recommendations from the RS.

Simplest example: the RS tells you the most popular items for other users until time t.

RS are now quite standard on the Web

Ecommerce sites: try to optimize revenue by getting users to buy more.

News and media sites: try to keep the users interacting with content, because it generates (advertising) revenue.

Users would get more utility if they receive recommendations they find valuable and trustworthy.

The formal study of RS

The study of RS is growing, from many perspectives:

- (Machine learning) algorithms - matrix approximation
- (Revenue) optimization - user metrics
- Controlled experiments

We study the effect of RS on markets.

Questions that interest us

Does the purchase of items have a steady-state distribution if an RS is influencing users?

(Under what conditions) do RS affect the purchase of items? How?

Is it possible for users to alter the popularities of items?

A model

A user may buy an item multiple times - e.g., a restaurant, batteries ...

<u>Prior purchases</u>: Each user has a list of prior purchases before the RS commences operation.

<u>Buying rate</u>: Each user makes purchases at a <u>frequency</u> f_u assume that these frequencies $\boldsymbol{f}$ add up over users to 1.

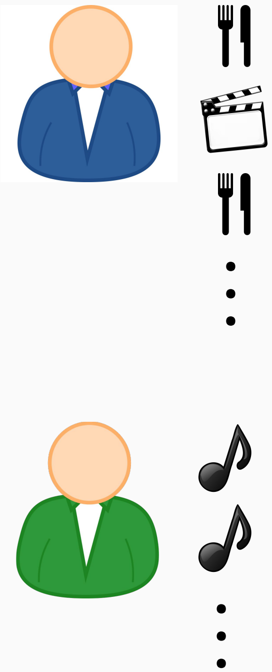

Model, continued

User's preference distribution: User *u* has a distribution $\boldsymbol{B}_u$ over the items reflecting its personal preferences. Absent the RS, *u*'s purchases are drawn from this distribution.

A probability α_u with which *u* consults the RS to determine what item to buy. With prob. 1-α_u it draws from $\boldsymbol{B}_u$ to make its purchase.

How does our RS work?

We have a directed, weighted, simple graph *G* connecting the users.

The weight w_{uv} denotes how much *u* “trusts” *v*.

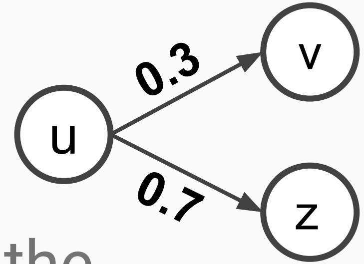

When *u* consults the RS, it picks a *v* according to the weights, then buys an item from the multiset of items bought by *v*.

One last modeling detail

We allow for recency when invoking the RS: we have a probability distribution ***h*** on the past purchases of the recommender *v* for a user *u*.

Examples:

Exponentially decaying on past purchases.

Only the last 10 purchases of *v*.

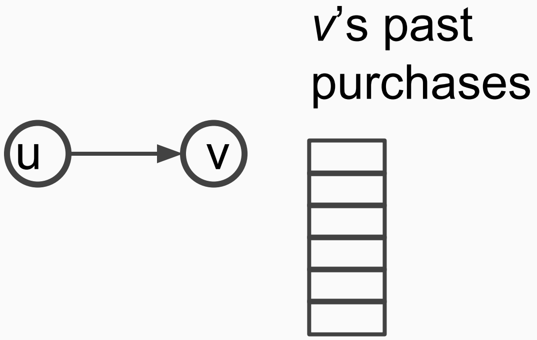

Now run this system - what happens?

Does the purchase rate of items have a steady-state distribution?

What is the effect of the RS on item popularities?

Are some users more influential than others?

A natural condition: The past fades away

If ***h*** **eventually** puts zero probability on the item purchased at any fixed time $<t$, we say the past fades away.

E.g., if uniform on all past purchases, or truncated beyond recent history.

Thus ***h*** will not put a fixed probability on a single purchase from long ago.

Main theorem: focus on a single product p^*

Let $\boldsymbol{A}$ = diag(... α_u...)

$\boldsymbol{b}$ be the vector of user preferences for p^*

$\boldsymbol{M}$ = [... w_{vu} ...], the weighted transposed adj matrix of G

$\boldsymbol{x}^t$ be the per-user fraction of purchases of p^* at time t

$\mathrm{Lim}_{t\to\infty} E[\boldsymbol{x}^t] = [(\boldsymbol{I}-\boldsymbol{AM})^{-1}(\boldsymbol{I}-\boldsymbol{A})]\boldsymbol{b} = \boldsymbol{Lb} = \boldsymbol{x}^\infty$

Influence of recommendation system

Influence and distortion

The relative contribution of each u to the overall sales distribution is an influence vector over users: $\boldsymbol{\gamma} = \boldsymbol{f}\boldsymbol{L}$

Define the market distortion as $\boldsymbol{\Delta} = (\boldsymbol{\gamma}\boldsymbol{b})/(\boldsymbol{f}\boldsymbol{b})$: the ratio of the market share of p^* with/without the RS.

Influence and some known precedents

If all $\alpha_u = \alpha$, then γ turns out to be exactly the vector of personalized pageranks [HKJ2003]

Consider a coalitional game where each user either joins the coalition for p^* or sits out. Then the Shapley value of user u is $\gamma_u b_u$

Convergence and computation

Convergence is fast for uniform recommendations from history.

To compute $\boldsymbol{\gamma}=\boldsymbol{fL}$ efficiently, expand the infinite summation $(\boldsymbol{I}\text{-}\boldsymbol{AM})^{-1}(\boldsymbol{I}\text{-}\boldsymbol{A})$ and truncate the series.

Experiments

How do influence and distortion play out in real social networks and RS?

Public datasets drawn from Google+, Twitter, Slashdot, Yelp and Facebook.

Set all α to be 0.2, which is abnormally large, to see how much we can distort the market.

Convergence

$||\boldsymbol{x}^t - \boldsymbol{x}^\infty||$ does appear to converge when history is forgotten by uniformly sampling from all past purchases.

It does not converge when history is forgotten through recency (settles at a fixed value), but $||E[\boldsymbol{x}^t] - \boldsymbol{x}^\infty||$ does appear to converge, as predicted by the theorem.

Convergence

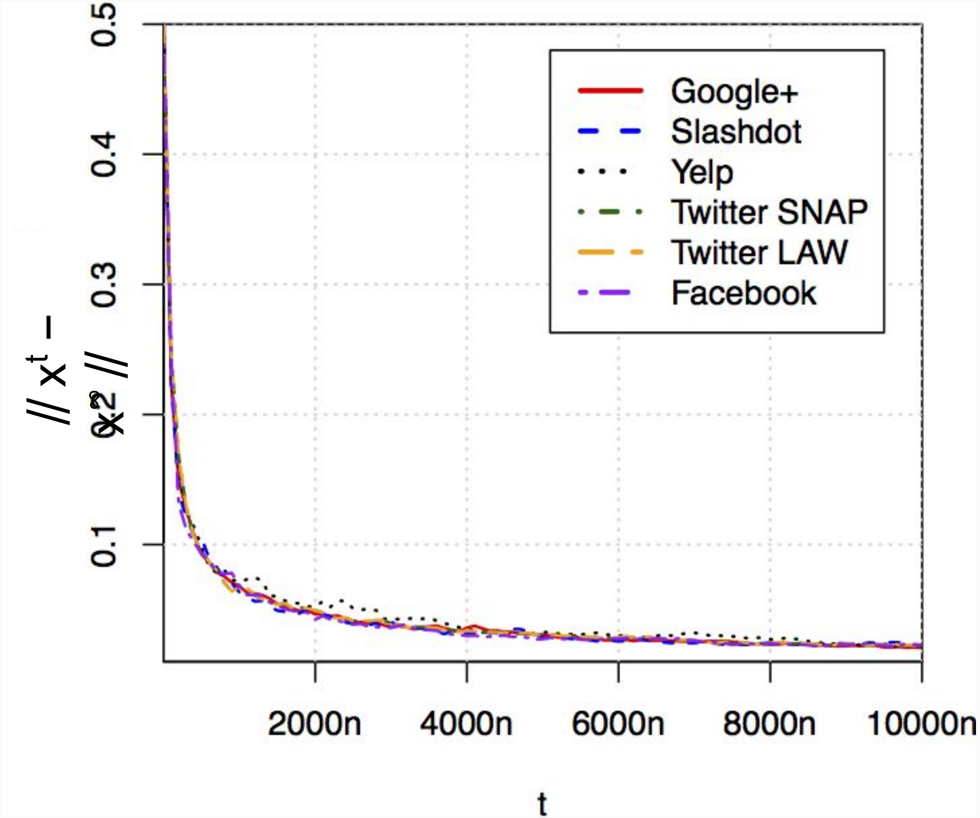

$||\boldsymbol{x}^t - \boldsymbol{x}^\infty||$ does converge when history is forgotten by uniformly sampling from all past purchases.

Convergence

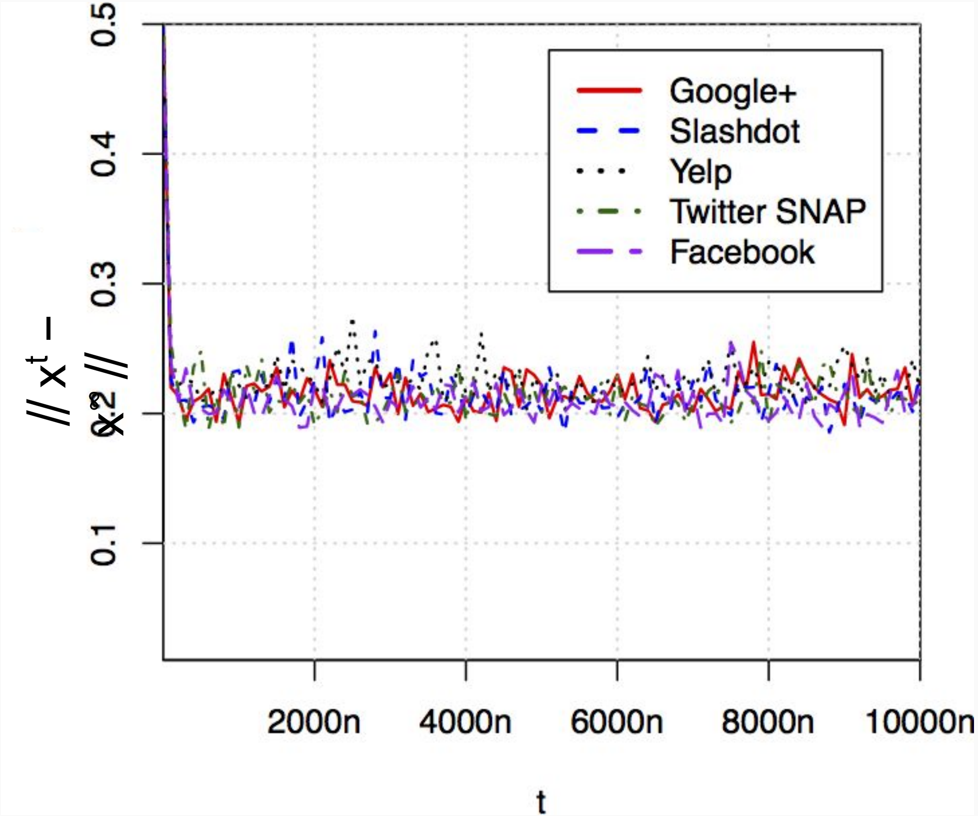

$||x^t - x^\infty||$ does not converge when history is forgotten through recency (however, $||E[x^t] - x^\infty||$ does).

In figure: each user remembers only his 100 most recent purchases.

Does the RS distort the market?

“Real” social networks dampen influence pretty heavily, with **Δ** always measured in the range 1 ± 0.002

On the other hand, planting an oligarchy (where everyone follows a few superstars) results in high **Δ**

Influence

In these “real” social networks, it was difficult to build large influence.

This hold even under “nonlinear” experiments where items were recommended with heavier probability than sampling.

Influence

Relative contribution of the most influential user (i.e. max γ over all nodes)

The same, expressed in “equivalent buying users” (i.e. max $n\gamma$ over all nodes)

Summary

Fairly general model of RS with a view to studying market influence.

Closed form for equilibrium, influence and market distortion.

Experiments suggest that “real” social networks dampen the influence of RS.