On the impact of social cost in opinion dynamics

Panagiotis Liakos **Katia Papakonstantinopoulou**

University of Athens

Algorithmic Game Theory Athens
NTUA, July 20th, 2016

Formation of opinions in a social context

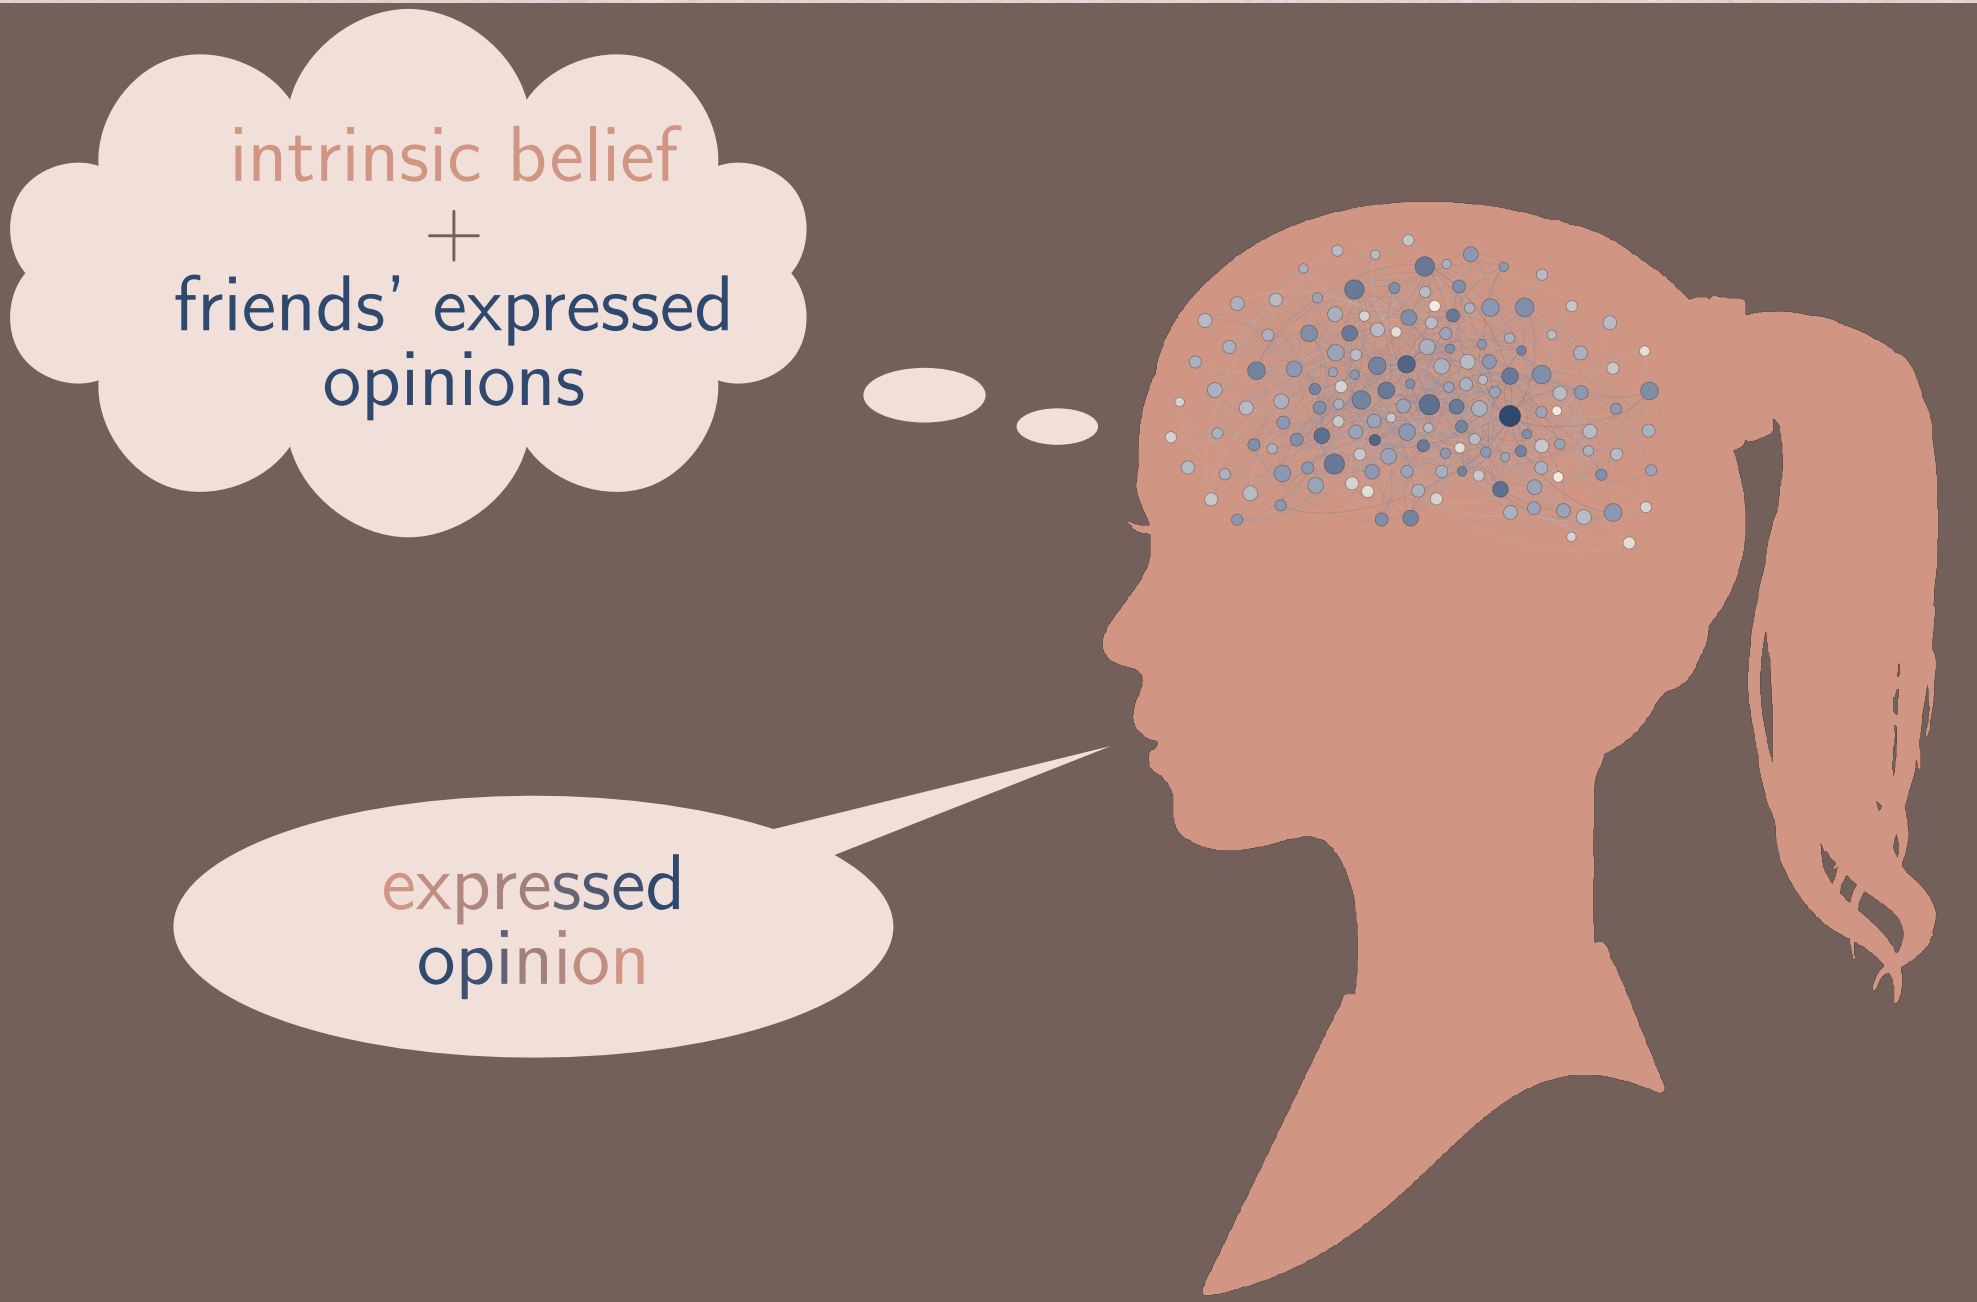

Motivation

exponential growth of online social networks

$\Downarrow$

ever-increasing amount of social activity information available

$\Downarrow$

ability to analyze user behavior and interpret sociological phenomena at a large scale [AKM08]

$\Downarrow$

Investigating game theoretic models of networks against real data

We consider the phenomenon of opinion formation under social influence. Given a network dataset, we want to be able to:

- verify the existence of influence among users
- build a model that describes user behavior in the network.

Contribution

Our contributions:

1. We analyze user activity in digg and verify that social interaction results in influence on opinions among the participants.
2. We initialize a sociological model using real data. Based on the Game Theory framework, we experimentally show that the *repeated averaging* process results to *Nash equilibria* which are illustrative of how users really behave.

What is digg?

digg is a news aggregator with a curated front page, aiming to select stories specifically for the Internet audience such as science, trending political issues, and viral Internet issues.

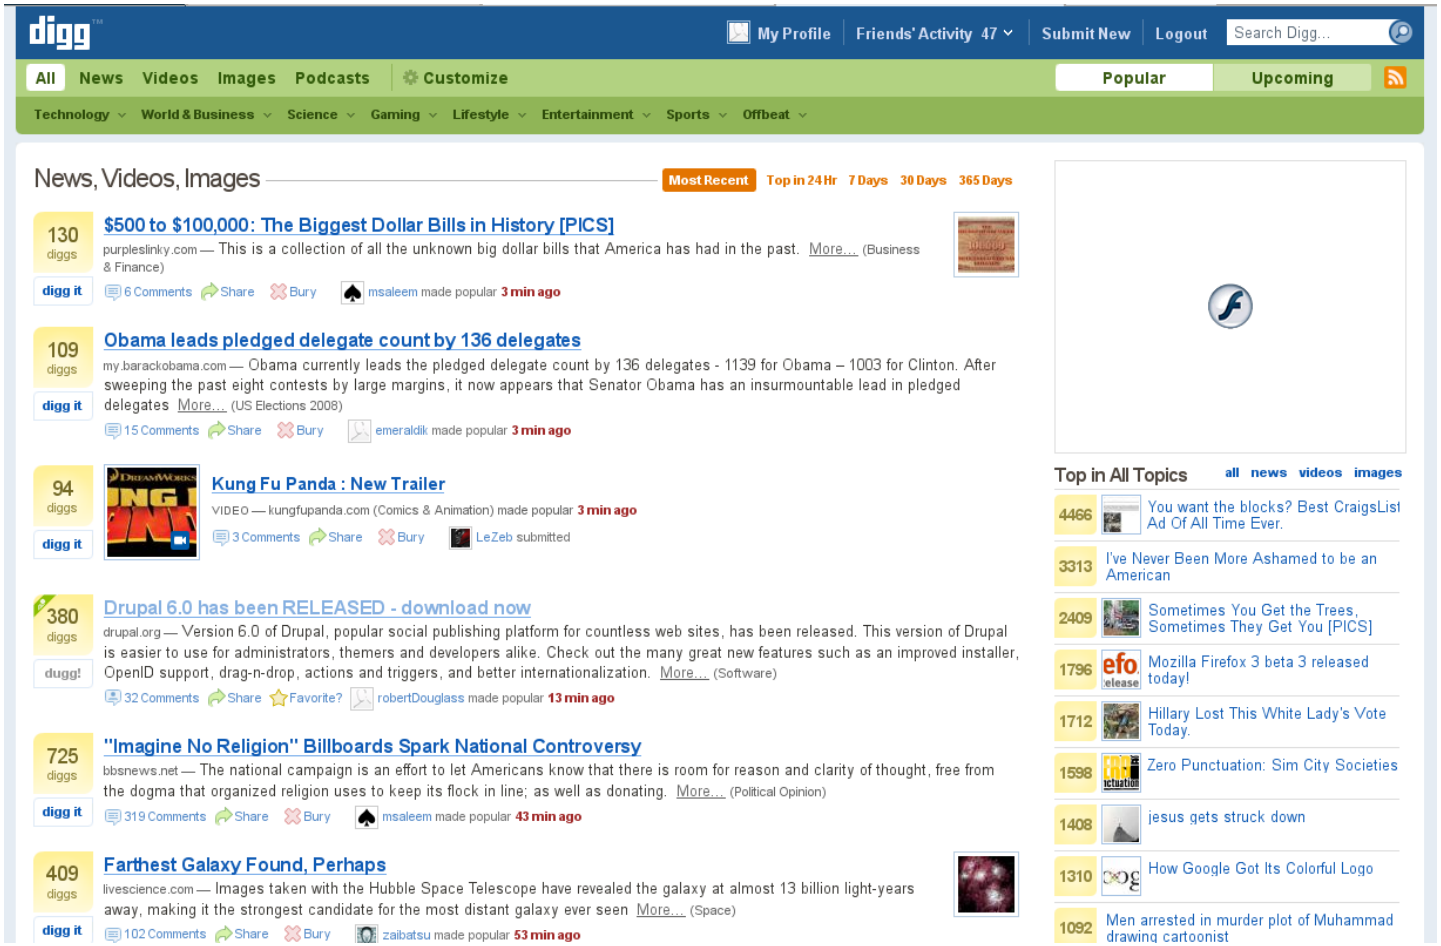

digg lets you:

- Submit stories.
- Digg (give a thumbs-up/positive vote to) a story you want other people to see.
- Follow users you consider interesting to get informed about their diggs in your news feed.

Why digg?

The digg dataset is appropriate for our study because:

- digg was very popular at the time the dataset was collected [LGS12]
- digging a story has a sense of opinion expression and an urge to influence
- both diggs and follower links are timestamped

The model: Basic notions

We use:

a variation of the **DeGroot** model due to Friedkin and Johnsen [FJ90] and the corresponding *game* of [BKO11].

Each user i maintains:

An *intrinsic* belief s_i
Remains constant

An *expressed* opinion z_i
Updated iteratively through averaging

The cost a user suffers emanates from:

Suppressing her intrinsic belief

Disagreeing with her friends

The model: Repeated Averaging

Repeated Averaging

At each time step user i updates z_i to *minimize* her cost:

$$z_i = \frac{s_i + \sum_{j \in N(i)} w_{ij} z_j}{1 + \sum_{j \in N(i)} w_{ij}}$$

$N(i)$: the set of nodes that i follows

w_{ij}: the *strength* of the influence of j on i

The averaging process terminates when z converges to the unique Nash equilibrium, where the social cost is minimized.

The model: Determining the influence strength

Our intuition:

The influence of j on i regarding a specific matter depends on:

The impact a_{ij} of j on i	The expertise b_j of j
Does i respect j's opinion in general?	Is j authoritative on this matter?

We define:

$$w_{ij} = a_{ij} b_j$$

Top-20 Cascade Patterns							
1		6	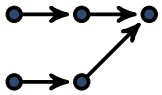	11	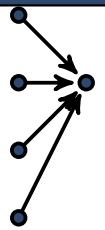	16	
2		7	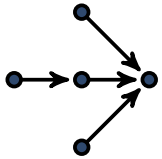	12		17	
3		8	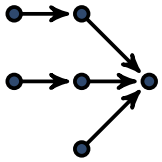	13	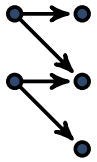	18	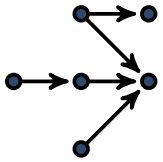
4	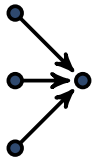	9	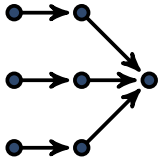	14	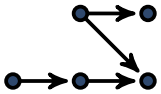	19	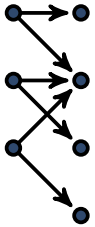
5	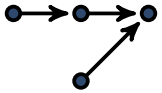	10		15	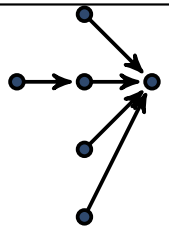	20	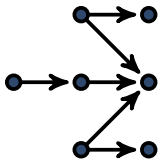

Empirical Analysis of digg (cont.)

the total *infections* were less than the initial *seeders* in more than 92% of the stories of digg

Empirical Analysis of digg (cont.)

the total *infections* were less than the initial *seeders* in more than 92% of the stories of digg

a few users caused many cascades while most were unable to cause any

Empirical Analysis of digg (cont.)

the total *infections* were less than the initial *seeders* in more than 92% of the stories of digg

a few users caused many cascades while most were unable to cause any

even the most *authoritative* users were not effective in all stories

Our Experiments and Assumptions

- We performed repeated averaging in our model for *stories* of **digg** until the opinions, expressed by votes, converged to the unique Nash equilibrium.
- We compared against predictions obtained using a *Neural Network* classifier.

Model initialization assumptions

Intrinsic belief s_i

$$s_i = \begin{cases} 1 & \text{if } \boldsymbol{i} \text{ voted a story before } \textit{any} \text{ user she follows} \\ 0 & \text{otherwise} \end{cases}$$

Influential strength w_{ij}

Two variants:

i) $a_{ij} = b_j = 1 \Rightarrow w_{ij} = 1$

ii) $a_{ij} = \frac{\#\text{ times } i \text{ is influenced by } j}{\#\text{ votes of } j}$

$b_j = \frac{\#\text{ users influenced by } j \text{ in this story}}{\#\text{ followers of } j}$

Results

The fraction of *predicted* votes that were actually *casted* against the fraction of *casted* votes that are *predicted*:

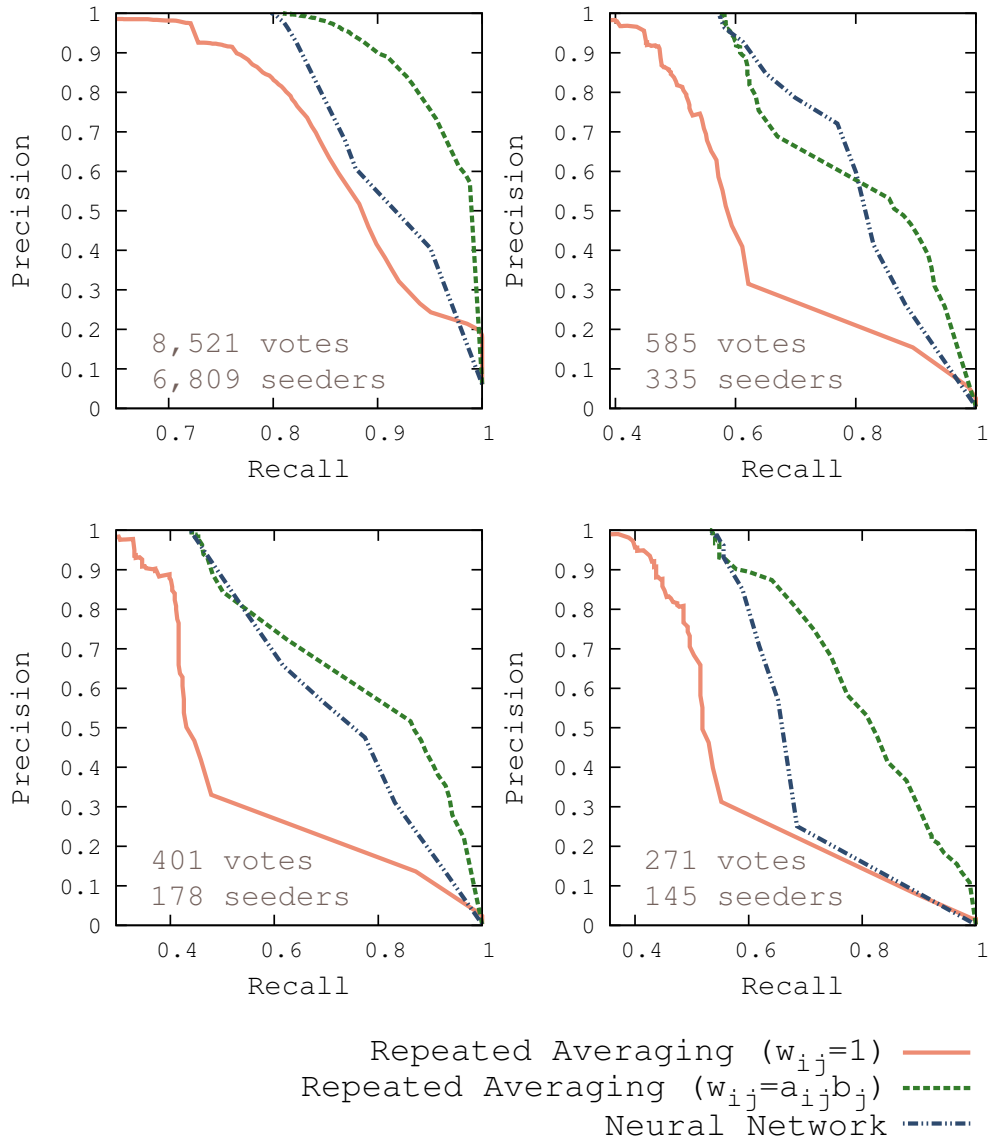

Results

The fraction of *predicted* votes that were actually *casted* against the fraction of *casted* votes that are *predicted*:

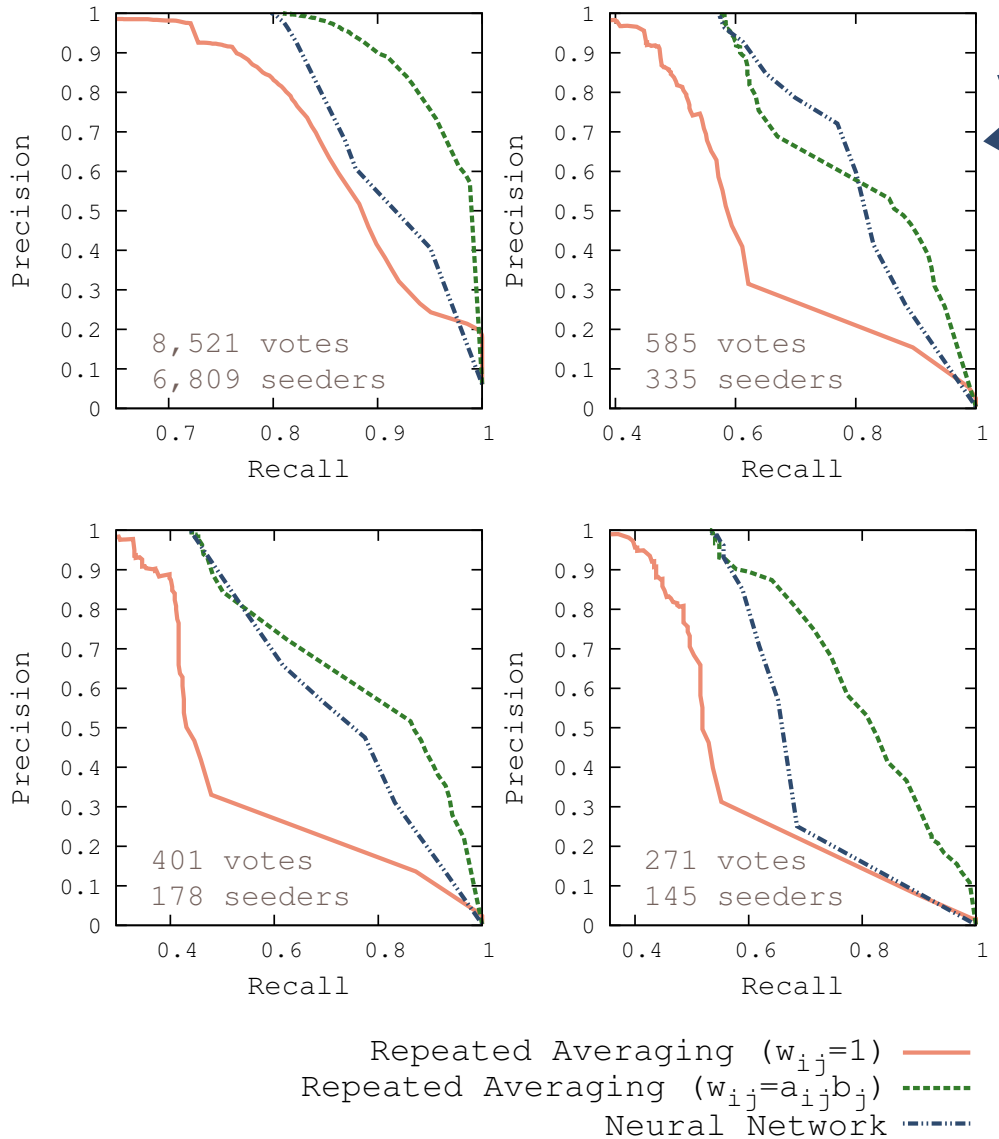

our 2nd variant closely mimics the original social activity

Results

The fraction of *predicted* votes that were actually *casted* against the fraction of *casted* votes that are *predicted*:

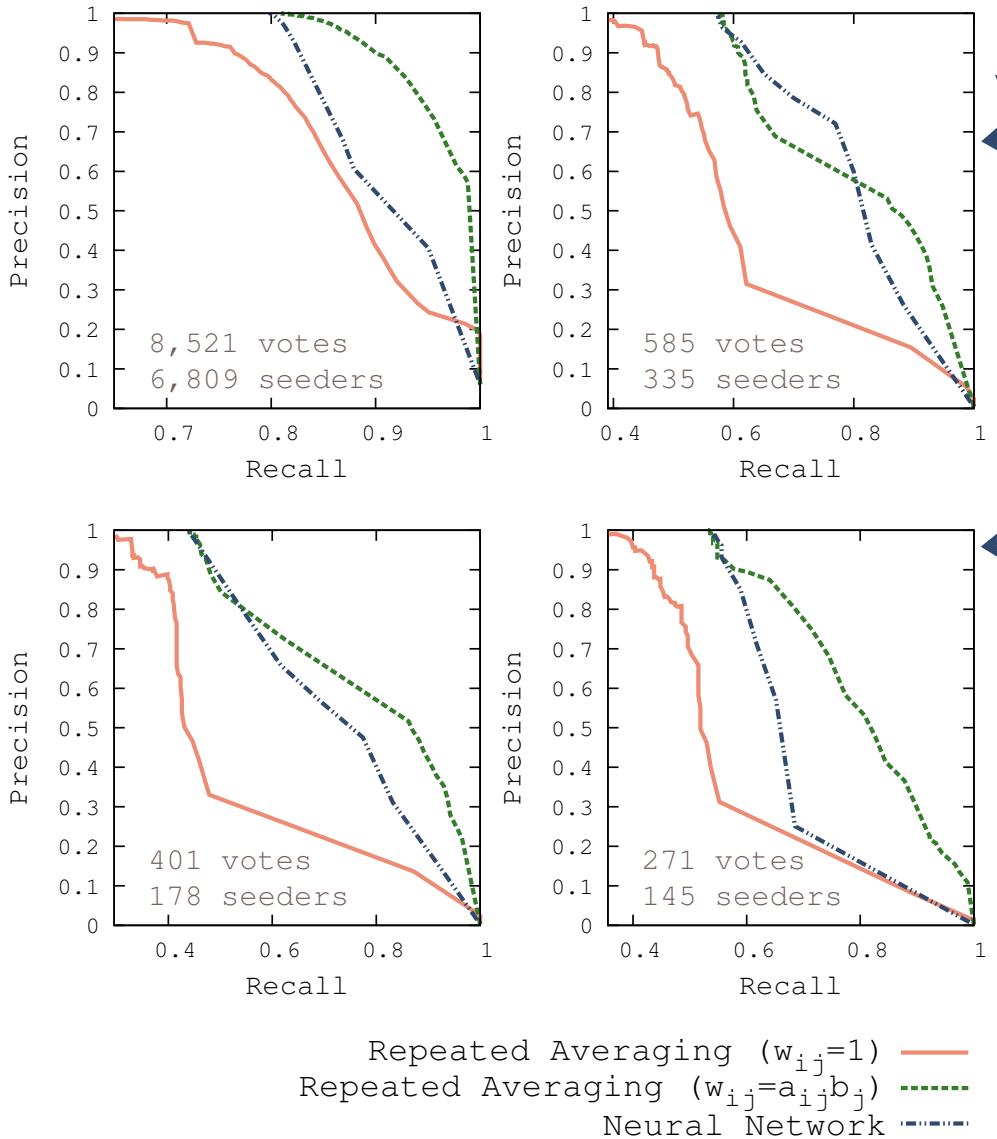

our 2nd variant closely mimics the original social activity

the simplistic variant behaves poorly as expected

Results

The fraction of *predicted* votes that were actually *casted* against the fraction of *casted* votes that are *predicted*:

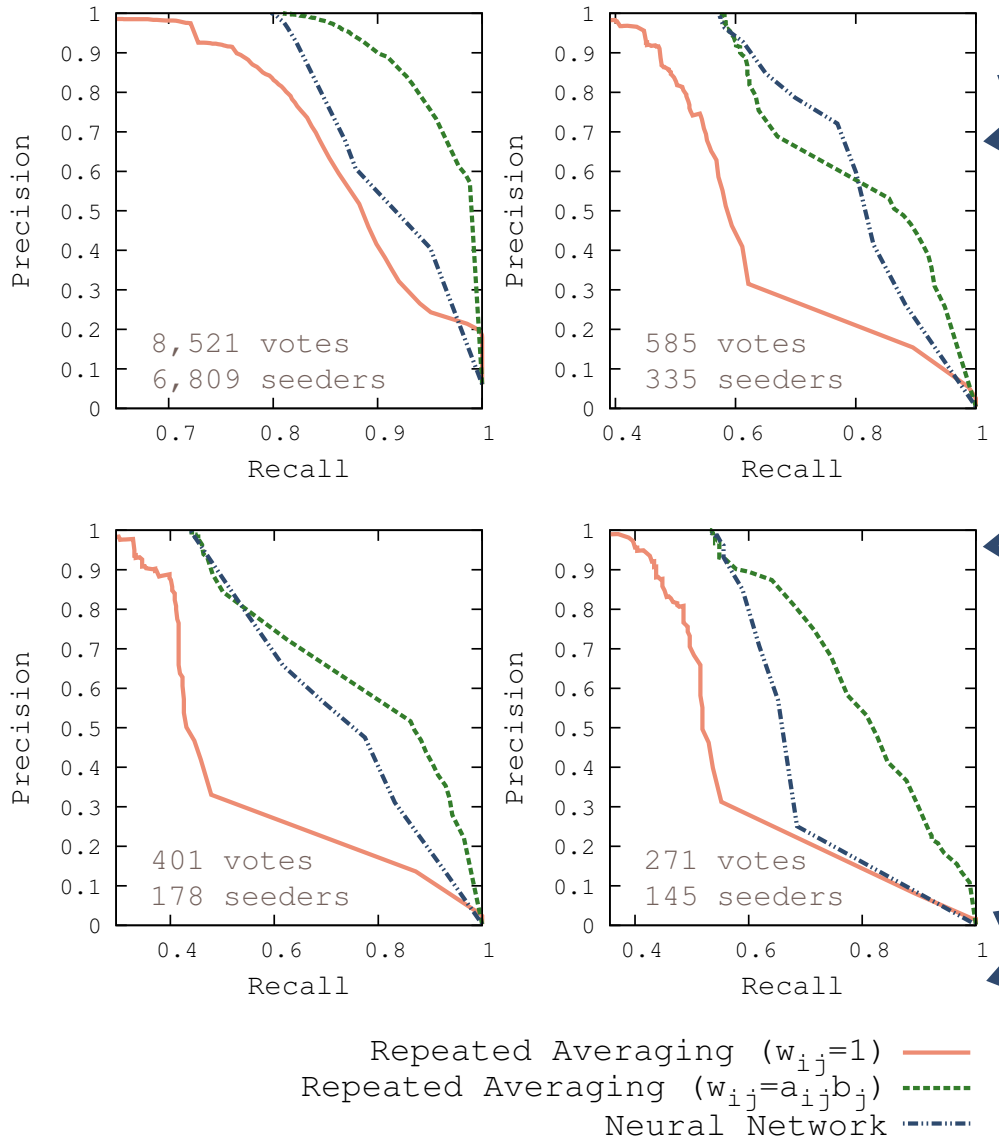

our 2nd variant closely mimics the original social activity

the simplistic variant behaves poorly as expected

we consistently out-perform the Neural Network classifier

Results

The fraction of *predicted* votes that were actually *casted* against the fraction of *casted* votes that are *predicted*:

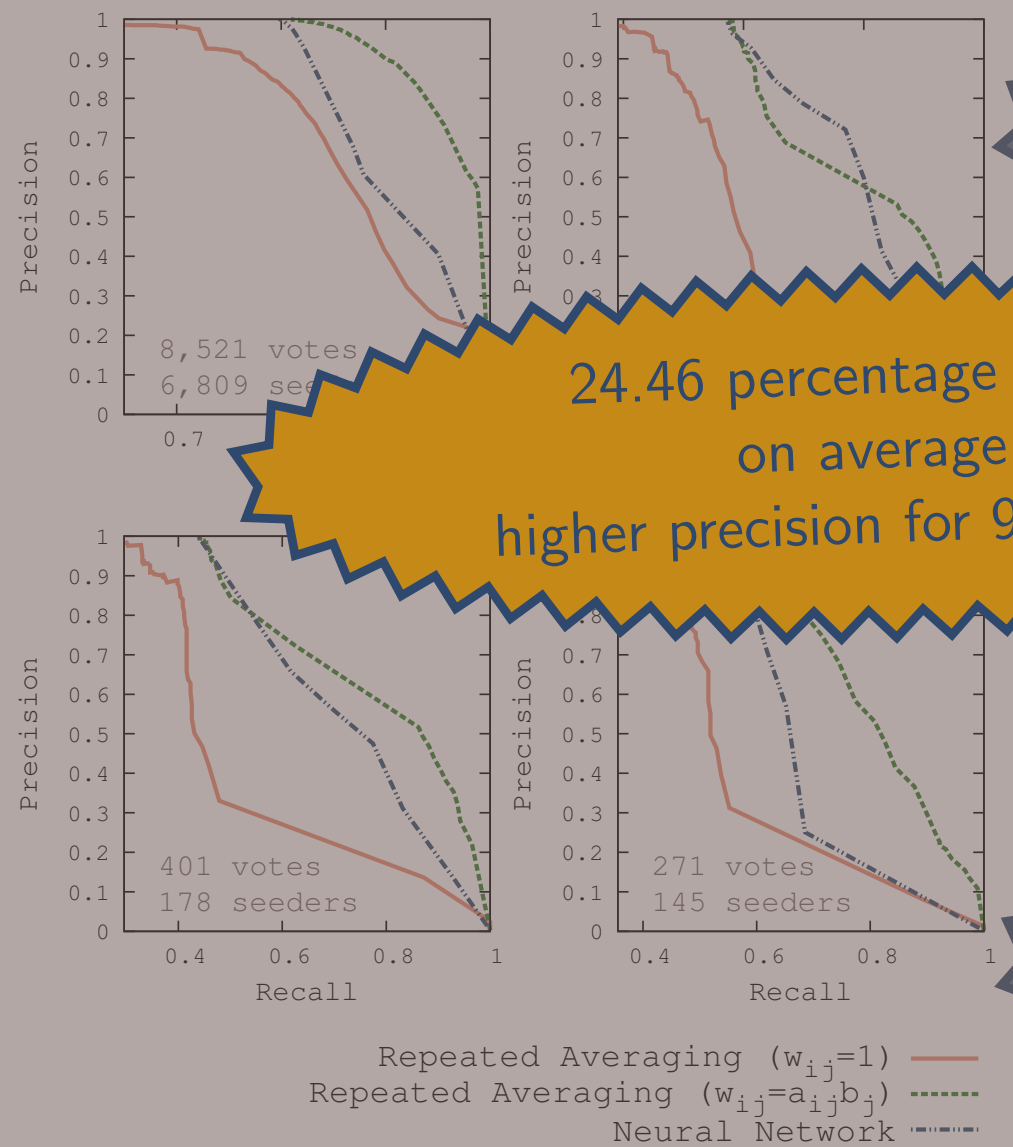

our 2nd variant closely mimics the original social activity

24.46 percentage points, on average, higher precision for 90% recall!

istic
behaves poorly as expected

we consistently out-perform the Neural Network classifier

Conclusion

The fraction of *predicted* votes that were actually *casted* against the fraction of *casted* votes that are *predicted*:

our 2nd variant closely mimics the original social activity

The repeated averaging process, combined with proper influence weights, results to Nash equilibria which are illustrative of how users really behave.

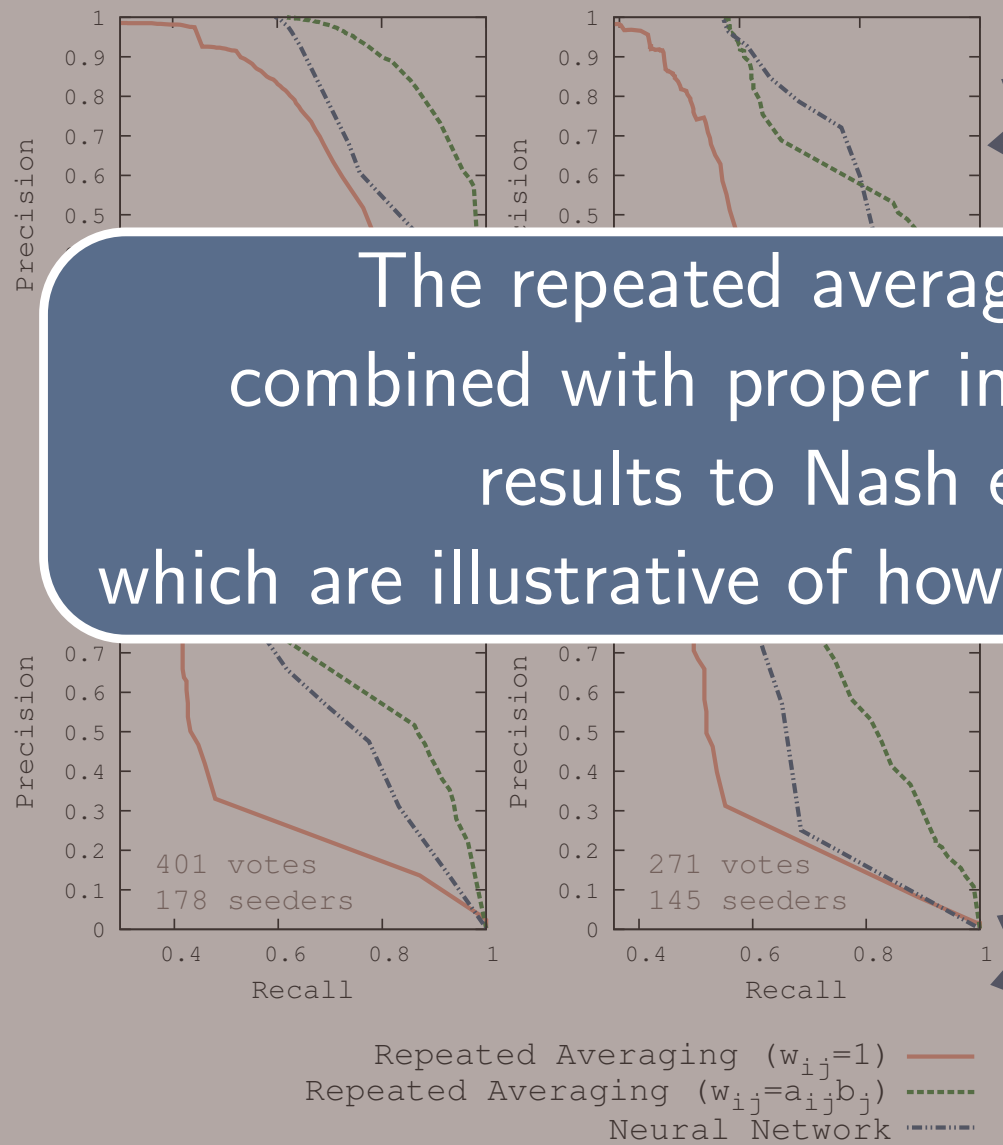

we consistently outperform the Neural Network classifier

References

[AKM08] Aris Anagnostopoulos, Ravi Kumar, and Mohammad Mahdian, *Influence and correlation in social networks*, SIGKDD (Las Vegas, Nevada, USA), 2008, pp. 7–15.

[BKO11] David Bindel, Jon M. Kleinberg, and Sigal Oren, *How bad is forming your own opinion?*, FOCS, 2011, pp. 57–66.

[FJ90] N.E. Friedkin and E.C. Johnsen, *Social influence and opinions*, Journal of Mathematical Sociology **15** (1990), no. 3-4, 193–206.

[LGS12] Kristina Lerman, Rumi Ghosh, and Tawan Surachawala, *Social contagion: An empirical study of information spread on digg and twitter follower graphs*, arXiv preprint arXiv:1202.3162 (2012).

Future Directions

- We can estimate the price of anarchy in such a network, by comparing the social cost of the Nash equilibrium we computed here with the cost of the optimal setting of this network.
- We can investigate approaches that may reduce the cost of such networks.

thank you!

for further details visit:
http://hive.di.uoa.gr/network-analysis/

or email me at: katia@di.uoa.gr