Opinion Diffusion as Aggregation

Umberto Grandi

IRIT – University of Toulouse

18 May 2017

[Based on joint work with Sirin Botan, Markus Brill, Edith Elkind, Ulle Endriss, Emiliano Lorini, Arianna Novaro and Laurent Perrussel]

Have you ever had an opinion?

What should
we do with Putin?

Have you ever had an opinion?

What should
we do with Putin?

Relax Angela!
Let's take a selfie first

Have you ever had an opinion?

What should
we do with Putin?

Relax Angela!
Let's take a selfie first

Have you ever had an opinion?

What should
we do with Putin?

Remember you are
a Nobel laureate
for peace

Attack
then think

Leave the
issue to the
next president

Relax Angela!
Let's take a selfie first

Have you ever had an opinion?

What should
we do with Putin?

Remember you are
a Nobel laureate
for peace

Attack
then think

Leave the
issue to the
next president

Relax Angela!
Let's take a selfie first

Opinion diffusion in the literature

Estabilished models in mathematical social sciences/social network analysis:

1. Discrete models
   - First studies by Granovetter and Schelling (1978)
   - Threshold models: 0/1 or yes/no opinions, updated if the proportion of neighbours with the opposite opinion raises above a certain threshold
   - Voter models (Holley and Ligget, 1975, Clifford and Sudbury, 1973): a random individual takes the opinion of a random neighbour
2. Continuous models
   - First studies by De Groot (1974) and Lehrer-Wagner (1981)
   - Opinions in [0,1] on a weighted network. Individuals take the weighted average of the opinions of their neighbours
   - First recent study involving logical constraints by Friedkin et al. (2016)

Opinion diffusion in the literature

Estabilished models in mathematical social sciences/social network analysis:

1. Discrete models
   - First studies by Granovetter and Schelling (1978)
   - Threshold models: 0/1 or yes/no opinions, updated if the proportion of neighbours with the opposite opinion raises above a certain threshold
   - Voter models (Holley and Ligget, 1975, Clifford and Sudbury, 1973): a random individual takes the opinion of a random neighbour
2. Continuous models
   - First studies by De Groot (1974) and Lehrer-Wagner (1981)
   - Opinions in [0,1] on a weighted network. Individuals take the weighted average of the opinions of their neighbours
   - First recent study involving logical constraints by Friedkin et al. (2016)

Two critical features of this literature:

1. The representation of individual opinions is too simple
2. An obsession with reaching consensus?

Opinion diffusion as aggregation

Opinions can be more complex than single 0/1 views or parameters in [0,1]:

- Vectors of binary views (Grandi et al. 2015)
- Preferences as linear orders over candidates (Brill et al. 2016)
- Belief bases as sets of propositional formulas (Schwind et al. 2015)

How are individuals updating on complex opinions?
A simple idea is to look at the opinion of one's influencers and:

- Uses aggregation rules from judgement aggregation (constraints!)
- Uses aggregation rules from preference aggregation (transitivity!)
- Uses belief merging techniques

Outline

1. A quick summary of related work + diffusion as aggregation (done)
2. Some general but useful definitions
3. The simplest case: vectors of binary opinions
4. Preferences and pairwise updates
5. Let's give life to the network nodes!
6. Conclusions and perspectives

The architecture of a discrete time iterated diffusion process - Part I

In virtually all settings there are common features:

- A finite set of individuals $\mathcal{N} = \{1, \ldots, n\}$
- A network $E \subseteq \mathcal{N} \times \mathcal{N}$ (typically directed)
- Individual opinions (unspecified format for now) that we shall denote as B_i

Some further notation: $\mathit{Inf}(i) = \{j \mid (i,j) \in E\}$ is the set of influencers of individual i on E. Profile of opinions are $\boldsymbol{B} = (B_1, \ldots, B_n)$.

An aggregation function for individual opinion updates

Each individual $i \in \mathcal{N}$ is provided with a suitably defined F_i that merge a set of opinions into a single one. The updated opinion of i is $F_i(B_i, \boldsymbol{B}\restriction_{\mathit{Inf}(i)})$.

Examples: F_i is the majority rule, a belief merging operator...

The architecture of a discrete time iterated diffusion process - Part II

Opinion diffusion process

Let $turn : \mathbb{N} \to 2^{\mathcal{N}}$ *indicate at each point in time the set of agents updating. Let* B_i^t *be the opinion of agent* i *at time* $t \in \mathbb{N}$*, and:*

$$B_i^{t+1} = \begin{cases} F_i(B_i, \boldsymbol{B}^t\restriction_{Inf(i)}) & \textit{if } i \in turn(t) \\ B_i^t & \textit{otherwise.} \end{cases}$$

If $turn(t) = \mathcal{N}$ the process is synchronous, if $turn$ selects one individual (typically uniformly at random) the process is asynchronous.

Disclaimer: when opinions are complex - multi-issue or preferences - we will also specify at each point in time the issue on which the update is performed.

Termination

Two forms of termination of the iterative process can be investigated:

Asymptotic termination

A diffusion model asymptotically terminates *on a class of graphs* $\mathcal{E} \subseteq 2^{N^2}$ *if for each graph* $E \in \mathcal{E}$ *and for each initial profile of opinions* $\boldsymbol{B}^0$ *we have*

$$\lim_{t \to +\infty} \mathbb{P}[\boldsymbol{B}^{t+1} \neq \boldsymbol{B}^t] = 0.$$

Typically applied to asynchronous models.

Universal termination

A diffusion model universally terminates *on a class of graphs* $\mathcal{E}$ *if there does not exist an infinite sequence of effective updates (ie. such that* $\boldsymbol{B}^{t+1} \neq \boldsymbol{B}^t$*).*

Typically hard to guarantee.

Convergence

Call a profile $\boldsymbol{B}^t$ *stable* if $F_i(\boldsymbol{B}^t) = B_i^t$ for all i, and a *termination profile* for $\boldsymbol{B}^0$ any stable profile reachable from $\boldsymbol{B}^0$.

What happens when the process terminates?

- Diffusion converges to unique profile if termination profiles coincide
- Diffusion converges to consensus if termination profiles are unanimous
- Other notions are of course possible...

The simplest case:
Multiple binary issues

A known example revisited

An influence network between Ann, Bob and Jesse:

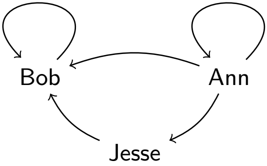

The three agents need to decide whether to approve the building of a swimming pool (first issue) and a tennis court (second issue) in the residence where they live. Here are their initial opinions and their evolution following propositional opinion diffusion with each agent syncronously using the majority rule:

Initial opinions	Profile $\boldsymbol{B}^1$	Profile $\boldsymbol{B}^2$
$B_A^0 = (0, 1)$	$B_A^1 = (0, 1)$	$B_A^2 = (0, 1)$
$B_B^0 = (0, 0)$	$B_B^1 = (0, 0)$	$B_B^2 = (0, 1)$
$B_J^0 = (1, 0)$	$B_J^1 = (0, 1)$	$B_J^2 = (0, 1)$

General termination result

A directed-acyclic graph (DAG) with loops is a directed graph that does not contain cycles involving more than one node.

Theorem

If F_i satisfies ballot-monotonicity for all i (see paper), then synchronous POD converges on the class of DAG with loops in at most diam$(E) + 1$ steps.

Proof. Start from the sources and propagate opinions.

Observations:

- The proof is a polynomial algorithm to compute the termination profile
- The theorem is not easy to strenghten: take the example of a circle
- The theorem works for any aggregator F_i, even those that do not treat issues independently

Summary of results

We show polynomial algorithms for the computing POD, and identify classes of graphs on which termination is guaranteed (only sufficient conditions):

Aggregation	Class of graphs	Time bound
Any aggregator	DAG with loops	$\textit{diam}(E) \times \textit{Time}(F)$
Unanimity rule	No loops, disjoint cycles, $\lvert \mathit{Inf}(i) \rvert > 1$ for one node	$O(n^2 m)$
Majority rule	Disjoint cycles $\lvert \mathit{Inf}(i) \rvert$ even on cycles	$O(n^2 m)$

Where $n = |\mathcal{N}|$ is the number of individuals, $m = |E|$ the number of arcs.

Follow-up work I

Many of the previous termination results are subsumed by the following paper:

Zoé Christoff and Davide Grossi. *Stability in Binary Opinion Diffusion*. Under submission.

Their result in a nutshell:

Definition - Winning and veto coalitions

$C \subset \mathcal{N}$ is a winning coalition for p and F if C can force acceptance of p by accepting it unanimously. $C \subset \mathcal{N}$ is a veto coalition for p and F if C can force rejection of p by rejecting it unanimously.

Theorem

Propositional opinion diffusion for independent and monotonic F terminates universally from profile $\boldsymbol{B}$ if and only if a specific condition relating winning and veto coalitions of F with profile $\boldsymbol{B}$ and network G holds.

Follow-up work II

What happens when issues are connected by an integrity constraint IC?
To avoid discursive dilemmas we could restrict updates to subsets of issues.

Definition - F-consistency

An opinion diffusion process is F-consistent if for all termination profiles $\boldsymbol{B}$ and all agents $i \in \mathcal{N}$, if $F(B_i, \boldsymbol{B}\restriction_{Inf(i)}) \in$ IC, then $B_i = F(B_i, \boldsymbol{B}\restriction_{Inf(i)})$.

Definition - Open structure

IC has open structure if any two models of IC at Hamming distance k are connected by a sequence of at most k models at distance $\leqslant 1$ from each other.

Theorem

Propositionwise diffusion is F-consistent if and only if IC has open structure.

Sirin Botan, Umberto Grandi and Laurent Perrussel. Propositionwise Opinion Diffusion with Constraints. In *Proceedings of EXPLORE-2017*, 2017.

Follow-up work III

What if an external agent want to manipulate the diffusion process?

Asynchronous majoritarian diffusion on undirected graphs and one single issue:

- Bribing individuals to change their opinions to obtain at least k zeros (or ones) at optimistic termination is NP-complete, $W[2]$-hard on number of vertices to bribe, $W[1]$-hard on tree-width of G.
- Removing edges to obtain at least k zeros (or ones) at optimistic termination is NP-complete ($+$ same parametrised results as above).
- Controlling the update sequence to get a specific profile at termination is NP-complete and W[1]-hard in tree-width of G.

Proofs. Reductions to Target Set Selection.

Robert Bredereck and Edith Elkind. Manipulating Opinion Diffusion in Social Networks. To appear in *Proceedings of IJCAI-2017*.

Preferences as linear orders over a set of candidates

The influence of a Condorcet cycle

An influence network with 4 agents and 3 alternatives. The preferences 1, 2, and 3 form a *Condorcet cycle*: the majority relation of their preferences is cyclic:

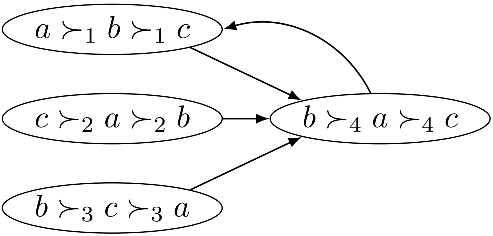

A possible branching of asynchronous pairwise preference diffusion (PPD):

- agent 4 updates on ab, moving to $a \succ_4 b \succ_4 c$
- no further updates possible: ac is no longer adjacent in $\succ_4$

A possible branching of synchronous PPD:

- agents 1 and 4 update repeatedly on pair ab
- an infinite update sequence starts

Convergence (or not) to consensus

A strongly connected network and a stable profile. PPD does not converge to consensus despite acyclic transitive majority:

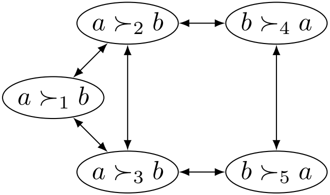

A simple cycle influence network and a preference profile with a cyclic majority relation. PPD converges to consensus:

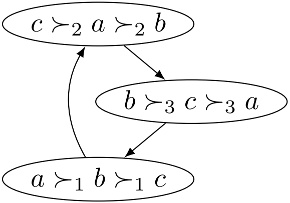

One interesting result and two open problems

Formalising the argument that mutual influence leads to aligned profiles:

Convergence to aligned profiles

If the sources of a DAG are aligned (single-peaked, single-crossing, Sen's restriction) then under mild conditions termination profiles are also aligned.

A number of problems are left open in pairwise opinion diffusion:

1. We show that asymptotic termination is guaranteed on all graphs (even cyclic ones) though under restrictive conditions (basically that no Condorcet cycle can ever occur). Can we relax this assumption?
2. We show that for less than three alternatives, a simple cycle can terminate on any linear order that does not conflict with a unanimous accepted pair in $\boldsymbol{B}^0$. How about the case when $|A| > 3$?

Let's give life to the network nodes!

An example: things get more complicated

Take a very simple network:

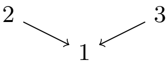

F_i is unanimity (change if all influencers agree). Consider two action profiles:

- $\boldsymbol{a}_0 = (\text{skip}, \text{skip}, \text{reveal}(p))$: agent 3 reveals her opinion
- $\boldsymbol{a}_1 = (\text{skip}, \text{hide}(p), \text{skip})$: agent 2 hides her opinion

Let a state be a profile of opinions (on one single issue p in this example), and a profile of visibilities. The two action profiles above generate a history:

$$\underbrace{((0,1,1),(1,1,0))}_{H_0} \xrightarrow[\boldsymbol{a}_0]{} \underbrace{((1,1,1),(1,1,1))}_{H_1} \xrightarrow[\boldsymbol{a}_1]{} \underbrace{((1,1,1),(1,0,1))}_{H_2}$$

Strategic disclosure of opinions on a social network

We define a simple notion of influence games:

- Individuals have private opinions on issues
- The actions at their disposal are reveal(p) and hide(p) for all issues p
- Each agent is endowed with an LTL goal

Example

Ann, Bob and Jesse have opinions (1,0,0) on a single issue p*:*

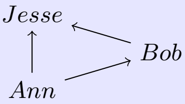

- *Ann's goal is* $\Diamond\Box\text{op}(Jesse, p)$
- *Memory-less winning strategy: play* reveal(p)

Outline of results

It is very hard to explore the existence of game-theoretic solution concepts restricting the structure of the networks and/or the goals: only limited results.

If we want to automate the search, then we need (lower bounds):

- LTL for memory-less strategies (PSPACE)
- ATL for perfect-recall winning strategy existence (EXPTIME)
- Graded strategy logic for Nash equilibrium existence (2EXPTIME)

A final observation: influence games are an interesting generalisation of iterated boolean games where propositional control is not exclusive!

Conclusions

To summarize:

- Opinion diffusion can be seen as iterated aggregation on a network
- The same idea can be applied to binary evaluations (eventually with constraints), preferences, or belief bases
- Problems to study: termination (asymptotic or universal), convergence (to consensus, to unanimity, to aligned profiles....)

Lots of open problems to be attacked. But (my) main interest is to go further:

Strategic aspects of collective decisions on networks

Recent papers showed that providing users with only local information (=looking at neighbouring nodes) generates processes that are resistant to a number of strategic actions by internal or external players (mostly on personalised recommender systems). How about voting rules?

References

This talk is based on the following papers:

Umberto Grandi, Emiliano Lorini and Laurent Perrussel. Propositional Opinion Diffusion. In *Proceedings of the 14th International Conference in Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.

Markus Brill, Edith Elkind, Ulle Endriss, and Umberto Grandi. Pairwise Diffusion of Preference Rankings in Social Networks. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

Sirin Botan, Umberto Grandi and Laurent Perrussel. Propositionwise Opinion Diffusion with Constraints. In *Proceedings of the 4th AAMAS Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE)*, 2017.

Umberto Grandi, Emiliano Lorini, Arianna Novaro and Laurent Perrussel. Strategic Disclosure of Opinions on a Social Network. In *Proceedings of the 16th International Conference in Autonomous Agents and Multiagent Systems (AAMAS)*, 2017.