Multi-Issue Opinion Diffusion under Constraints

Umberto Grandi

IRIT – University of Toulouse

5 June 2019

Joint work with Sirin Botan (Amsterdam)
and Laurent Perrussel (Toulouse)

Pre-vote discussions

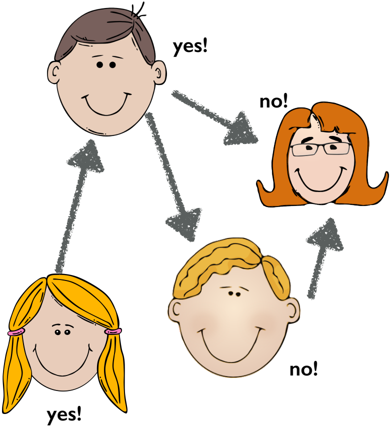

Mutual influence
(deliberation?)

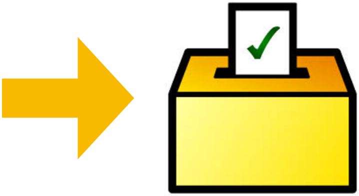

Vote!

One single issue (or multiple issues without constraints)

The model:

- n agents on a network E (directed/undirected)
- each agent has a 0/1 opinion
- the update is typically done by setting a threshold for each agent

One single issue (or multiple issues without constraints)

The model:

- n agents on a network E (directed/undirected)
- each agent has a 0/1 opinion
- the update is typically done by setting a threshold for each agent

Results known from the literature:

- Goles and Olivos (1980) showed that the process either terminates, or cycles with period 2
- Characterisations of profiles, networks, and aggregators that guarantee termination (previous work AAMAS-2015, Christoff and Grossi, 2017)
- Many papers characterising the termination profiles for the majority dynamics (including distinguished paper at IJCAI-2018)
- Strategic manipulation to maximise a given opinion under majority dynamics (Bredereck and Elkind, 2017)

Constrained collective choices I

Four individuals are deciding to build a skyscraper (S), a new road (R), or a hospital (H). Law says that if S and H are built then R also should be built.



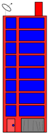

(Hosp and SkyS) implies Road

Voter 1: Y N N

Voter 2: N N Y

Voter 3: Y Y Y

Voter 4: N N N

Constrained collective choices I

Four individuals are deciding to build a skyscraper (S), a new road (R), or a hospital (H). Law says that if S and H are built then R also should be built.

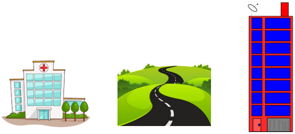

(Hosp and SkyS) implies Road

Voter 1: Y N N

Voter 2: N N Y

Voter 3: Y Y Y

Voter 4: N N N

What can happen:

- If voter 4 asks her influencers on 3 issues at the time then the update is blocked by an inconsistent issue-by-issue majority (Y N Y) (yes, this is an instance of the discursive dilemma).
- If voter 4 asks questions on a single issue to her influencers then the result can either be (Y N N) or (N N Y)

Outline

1. Aggregation-based opinion diffusion on multiple issues with constraints
2. Propositionwise updates and geodetic constraints
3. Cost of constraints and termination results
4. Conclusions and perspectives

Basic definitions

In virtually all settings there are common features:

- A finite set of individuals $\mathcal{N} = \{1, \dots, n\}$
- A finite set of issues or questions $\mathcal{I} = \{1, \dots, m\}$
- A directed graph $E \subseteq \mathcal{N} \times \mathcal{N}$ representing the trust network
- Individual opinions as vectors of yes/no answers $B \in \{0,1\}^{\mathcal{I}}$
- An integrity constraint $\mathrm{IC} \subseteq \{0,1\}^{\mathcal{I}}$

Basic definitions

In virtually all settings there are common features:

- A finite set of individuals $\mathcal{N} = \{1, \dots, n\}$
- A finite set of issues or questions $\mathcal{I} = \{1, \dots, m\}$
- A directed graph $E \subseteq \mathcal{N} \times \mathcal{N}$ representing the trust network
- Individual opinions as vectors of yes/no answers $B \in \{0,1\}^{\mathcal{I}}$
- An integrity constraint $\mathrm{IC} \subseteq \{0,1\}^{\mathcal{I}}$

A first example of the problems we consider:

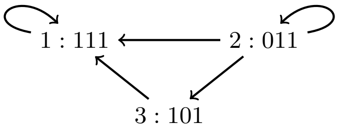

Diffusion as aggregation

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 $\boldsymbol{B} = (B_1, \dots, B_n)$.

An aggregation function for individual opinion updates

Each individual $i \in \mathcal{N}$ is provided with a suitably defined F_i that merge the set of opinions of its influencers into an aggregated view $F_i(\boldsymbol{B}\restriction_{\mathit{Inf}(i)})$.

Examples: F_i is the majority rule, a distance-based operator...examples can be found in the literature on judgment and binary aggregation (see Endriss, 2016)

We assume every F_i to be unanimous: if $B_i = B$ for all $i \in \mathcal{N}$ then $F(\boldsymbol{B}) = B$. No negative influence is possible in unanimous profiles.

Update simultaneously on all issues

When clear from the context F can represent an aggregation function or a profile of aggregation functions F_i, one for each agent.

Definition - Propositional opinion diffusion

Given network G and aggregators F, we call propositional opinion diffusion (POD) the following transformation function:

$$\begin{aligned} \mathrm{POD}_F(\boldsymbol{B}) = \{ & \boldsymbol{B}' \mid \exists M \subseteq \mathcal{N} \\ & \textit{s.t.} \quad B'_i = F_i(\boldsymbol{B}_{\textit{Inf}(i)}) \textit{ if } \mathrm{IC}\textit{-consistent and } i \in M \\ & \textit{and } B'_i = B_i \textit{ otherwise.} \} \end{aligned}$$

Update on subsets of issues

Definition - F-updates

Let F be an aggregation function, and let $(B\restriction_{\mathcal{I}\setminus S}, B'\restriction_S)$ be the opinion obtained from B with the opinions on the issues in S replaced by those in B'.

$$F\text{-UPD}(\boldsymbol{B}, i, S) = \begin{cases} (B_i\restriction_{\mathcal{I}\setminus S}, F_i(\boldsymbol{B}_{\mathit{Inf}(i)})\restriction_S) & \textit{if } \mathrm{IC}\textit{-consistent} \\ B_i & \textit{otherwise.} \end{cases}$$

Update on subsets of issues

Definition - F-updates

Let F be an aggregation function, and let $(B\restriction_{\mathcal{I}\setminus S}, B'\restriction_S)$ be the opinion obtained from B with the opinions on the issues in S replaced by those in B'.

$$F\text{-UPD}(\boldsymbol{B}, i, S) = \begin{cases} (B_i\restriction_{\mathcal{I}\setminus S}, F_i(\boldsymbol{B}_{\mathit{Inf}(i)})\restriction_S) & \textit{if } \mathrm{IC}\textit{-consistent} \\ B_i & \textit{otherwise.} \end{cases}$$

Definition - Propositionwise opinion diffusion

Given network G, aggregation functions F, and $1 \leqslant k \leqslant |\mathcal{I}|$, we call k-propositionwise opinion diffusion the following transformation function:

$$\begin{aligned} \mathrm{PWOD}_F^k(\boldsymbol{B}) = \{\boldsymbol{B}' \mid {} & \exists M \subseteq \mathcal{N}, S : M \to 2^{\mathcal{I}} \textit{ with } |S(i)| \leqslant k, \\ & \textit{s.t. } B'_i = F\text{-UPD}(\boldsymbol{B}, i, S(i)) \textit{ for } i \in M \\ & \textit{and } B'_i = B_i \textit{ otherwise.}\} \end{aligned}$$

Example

An influence network between four agents, with $\mathrm{IC} = (S \wedge H \rightarrow R)$:

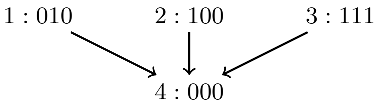

If F_4 the strict majority rule, then $F_4(B_1, B_2, B_3) = 110$. We have that:

- $\mathrm{POD}_F(\boldsymbol{B}) = \{\boldsymbol{B}\}$, we say that $\boldsymbol{B}$ is a termination profile for POD_F

Example

An influence network between four agents, with $\mathrm{IC} = (S \wedge H \to R)$:

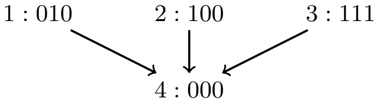

If F_4 the strict majority rule, then $F_4(B_1, B_2, B_3) = 110$. We have that:

- $\mathrm{POD}_F(\boldsymbol{B}) = \{\boldsymbol{B}\}$, we say that $\boldsymbol{B}$ is a termination profile for POD_F
- $\mathrm{PWOD}^1_F(\boldsymbol{B}) = \{(010, 100, 111, \mathbf{010}), (010, 100, 111, \mathbf{100}), \boldsymbol{B}\}$.

Example

An influence network between four agents, with $\mathrm{IC} = (S \wedge H \to R)$:

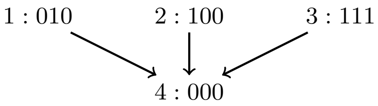

If F_4 the strict majority rule, then $F_4(B_1, B_2, B_3) = 110$. We have that:

- $\mathrm{POD}_F(\boldsymbol{B}) = \{\boldsymbol{B}\}$, we say that $\boldsymbol{B}$ is a termination profile for POD_F
- $\mathrm{PWOD}^1_F(\boldsymbol{B}) = \{(010, 100, 111, \mathbf{010}), (010, 100, 111, \mathbf{100}), \boldsymbol{B}\}$.
- $\mathrm{PWOD}^2_F(\boldsymbol{B}) = \mathrm{PWOD}^1_F(\boldsymbol{B})$

Problematic example

Let there be two issues and $\mathrm{IC} = p\,\mathrm{XOR}\,q = \{01, 10\}$. Consider the following:

$$1{:}\ 01 \longrightarrow 2{:}\ 10$$

Whatever the unanimous F:

- $\mathrm{POD}_F(\boldsymbol{B}) = \{\boldsymbol{B}, \boldsymbol{B}'\}$ where $B'_1 = B'_2 = (0, 1)$
- $\mathrm{PWOD}^1_F(\boldsymbol{B}) = \{\boldsymbol{B}\}$

Problematic example

Let there be two issues and $\mathrm{IC} = p \, \mathrm{XOR} \, q = \{01, 10\}$. Consider the following:

$$1\text{: } 01 \longrightarrow 2\text{: } 10$$

Whatever the unanimous F:

- $\mathrm{POD}_F(\boldsymbol{B}) = \{\boldsymbol{B}, \boldsymbol{B}'\}$ where $B'_1 = B'_2 = (0, 1)$
- $\mathrm{PWOD}^1_F(\boldsymbol{B}) = \{\boldsymbol{B}\}$

Question

Can we characterise the set of integrity constraints on which PWOD^k_F-reachability corresponds to POD_F-reachability?

Digression: k-geodetic integrity constraints

Observe that a constraint IC can be seen as a boolean function, and define:

Definition

The k-graph of IC is given by $\mathcal{G}^k_{\mathrm{IC}} = \langle \mathrm{IC}, E^k_{\mathrm{IC}} \rangle$, where:

1. *the set of nodes is the set of $B \in \mathrm{IC}$,*
2. *the set of edges E^k_{IC} is defined as follows: $(B, B') \in E^k_{\mathrm{IC}}$ iff $H(B, B') \leqslant k$, for any $B, B' \in \mathrm{IC}$.*

Where the Hamming distance $H(B, B')$ is the number of disagreements between two ballots B and B'.

Definition - Geodetic integrity constraints

An integrity constraint IC is k-geodetic if and only if for all B and B' in IC, at least one of the shortest paths from B to B' in $\mathcal{G}^k_{\top}$ is also a path of $\mathcal{G}^k_{\mathrm{IC}}$.

Examples I

- $\mathrm{IC} = \{(000), (001), (010), (100), (011), (111)\}$ is 2-geodetic but not 1-geodetic, as can be seen on $\mathcal{G}^1_{\mathrm{IC}}$:

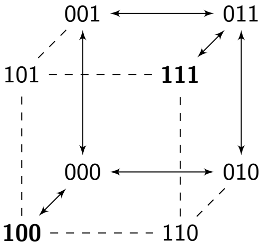

- Our running example $\mathrm{IC} = S \wedge H \to R = \{(000), (001), (010), (011), (100), (101), (111)\}$ is 1-geodetic, as only one model is missing.

Examples of 1-geodetic constraints

Preferences. Let $a > b$ be a set of binary questions for candidates $a, b, c \dots$. The constraints are that of transitivity, completeness and anti-symmetry. This set of constraints is 1-geodetic, since two distinct linear orders always differ on at least one adjacent pair.

Budget constraints. Enumerate all combinations of items that exceed a given budget. They are *negative formulas*, ie. one DNF representation only has negative literals: a sufficient condition for 1-geodeticity.

More examples of 1-geodetic boolean function/constraints in:

Ekin, Hammer, and Kogan. On Connected Boolean Functions. *Discrete Mathematics*, 1999.

Reachability result

Theorem

Let IC *be an integrity constraint. Any profile* $\boldsymbol{B}'$ *that is* POD_F*-reachable from an* IC*-consistent initial profile* $\boldsymbol{B}$ *is also* PWOD_F^k*-reachable from* $\boldsymbol{B}$ *if and only if* IC *is* k*-geodetic.*

Proof sketch.

$\Rightarrow$) If $\boldsymbol{B'}$ is reachable by updating all issues at the same time, then by k-geodeticity it is also reachable by updates on sets of issues of size k.

$\Leftarrow$) If IC is not k-geodetic there are two disconnected models. Construct a problematic example such as the one seen before (assumption of unanimity of F used here).

Theorem

Let IC *be a constraint over* m *issues and* $k < m$*. Checking whether* IC *is* k*-geodetic is* co-NP*-complete.*

Proof sketch.

For membership: Guess two models B and B' and check if all shortest paths connecting them start with a non-model of IC (this can be done in time polynomial in parameter k);

For completeness: use a result by Hegedus and Megiddo (1996) on classes of boolean functions that have the projection property.

Cost of constraints and termination

Question - Cost of constraints

Can we quantify the gain in terms of influence that is given by allowing for updates on k issues?

Answer: define the influence gap as the sum of the distances between every individual's opinion and the aggregated one of its infleuncers. We show that this figure for POD_F is larger than for PWOD_F^k and give precise bounds.

Cost of constraints and termination

Question - Cost of constraints

Can we quantify the gain in terms of influence that is given by allowing for updates on k issues?

Answer: define the influence gap as the sum of the distances between every individual's opinion and the aggregated one of its infleuncers. We show that this figure for POD_F is larger than for PWOD_F^k and give precise bounds.

Question - Termination

Can we find conditions on the graph and the aggregation functions that guarantee that the opinion diffusion will terminate?

Answer: similar findings as for single issue for what concerns complete graphs, DAGs, for arbitrary graphs we have to assume consistent aggregation of influencer's opinions. Open problem: can this last assumption be relaxed?

Conclusions

In this work:

- We started by viewing opinion diffusion as iterated aggregation on a network, adding integrity constraints
- We characterised the set of integrity constraints for which reachability when updating on all the issues implies propositionwise reachability (and assessed the gain in terms of Hamming distance)
- We showed initial results on the termination of such processes

Lots of open problems to be attacked:

- Can we relax the local consistency property? What is the class of constraints on which termination is guaranteed?
- Any relation between constraints and network structure to guarantee termination?
- Generalise to uncertain agents (yes-no-don't know)
- Strategic influence?

References

Previous work on the topic:

- S. Botan, UG and L. Perrussel. Multi-issue Opinion Diffusion Under Constraints. In *Proceedings of the 18th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2019.
- UG, E. Lorini, A. Novaro and L. Perrussel. Strategic Disclosure of Opinions on a Social Network. In *Proceedings of the 16th International Conference in Autonomous Agents and Multiagent Systems (AAMAS)*, 2017.
- M. Brill, E. Elkind, U. Endriss, and UG. Pairwise Diffusion of Preference Rankings in Social Networks. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- UG, E. Lorini and L. Perrussel. Propositional Opinion Diffusion. In *Proceedings of the 14th International Conference in Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.

Termination of POD and PWOD

Given a transformation function POD_F or PWOD_F^k, we can consider:

- **Asynchronous** opinion diffusion when only one agent at the time updates
- **Synchronous** opinion diffusion when all agents at the same time update

Basic definitions of iterative diffusion processes

Given a transformation function POD_F or PWOD_F^k, we can consider:

- **Asynchronous** opinion diffusion when only one agent at the time updates
- **Synchronous** opinion diffusion when all agents at the same time update

Two termination notions are possible:

- **Universal termination**: there exists no sequence of *effective* updates (ie when $\boldsymbol{B}_{t+1} \neq \boldsymbol{B}_t$)
- **Asymptotic termination**: from any IC-consistent profile there exists a sequence of updates to reach a termination profile

Universal termination

Ballot-Monotonicity: for all profiles $\boldsymbol{B} = (B_1, \ldots, B_n)$, if $F(\boldsymbol{B}) = B^*$ then for any $1 \leqslant i \leqslant n$ we have that $F(\boldsymbol{B}_{-i}, B^*) = B^*$.

Theorem

Let G be the complete graph. Synchronous POD_F terminates universally, and asynchronous POD_F terminates universally if F is ballot-monotonic.

Universal termination

Ballot-Monotonicity: for all profiles $\boldsymbol{B} = (B_1, \ldots, B_n)$, if $F(\boldsymbol{B}) = B^*$ then for any $1 \leqslant i \leqslant n$ we have that $F(\boldsymbol{B}_{-i}, B^*) = B^*$.

Theorem

Let G be the complete graph. Synchronous POD_F terminates universally, and asynchronous POD_F terminates universally if F is ballot-monotonic.

Monotonicity: for any $j \in \mathcal{I}$ and any profiles $\boldsymbol{B}, \boldsymbol{B}'$, if $B_i(j){=}1$ entails $B'_i(j){=}1$ for all $i \in \mathcal{N}$, and for some $s \in \mathcal{N}$ we have that $B_s(j){=}0$ and $B'_s(j){=}1$, then $F(\boldsymbol{B})(j){=}1$ entails $F(\boldsymbol{B}')(j){=}1$

Theorem

If G is the complete graph and F is monotonic, then both synchronous and asynchronous PWOD_F^k terminate universally.

Termination of asynchronous processes

A well-known construction generalises to k-geodetic integrity constraints.

Definition

A pair $(\boldsymbol{B}^0, G)$*, where* $\boldsymbol{B}^0$ *is a profile and* G *a network, has the local* IC*-consistency property if for all profiles* $\boldsymbol{B}$ *reachable from* $\boldsymbol{B}^0$ *and each* $i \in \mathcal{N}$ *we have that* $F(\boldsymbol{B}_{\mathsf{Inf}(i)})$ *is* IC*-consistent.*

Theorem

If $(\boldsymbol{B}^0, G)$ *satisfies the local IC-consistency property, then asynchronous* POD_F *and* PWOD_F^k *terminate asymptotically.*

Proof sketch

Fix an ordering of the issues. For each issue perform two following rounds:

- First round of updates: all individuals who disagree with their influencers and have opinion 0 update their opinion to 1
- Second round of updates: all individuals who disagree with their influencers and have opinion 1 update their opinion to 0