Trends in Computational Social Choice

9

Cite as: Umberto Grandi. Social Choice and Social Networks. In Ulle Endriss (editor), *Trends in Computational Social Choice*, chapter 9, pages 169–184. AI Access, 2017.

http://www.illc.uva.nl/COST-IC1205/Book/

CHAPTER 9

Social Choice and Social Networks

Umberto Grandi

9.1 Introduction

Individuals do not typically reason in isolation when confronted with collective-decision making, but rather take into consideration the preferences of like-minded individuals and engage in strategic activities such as influence, persuation and information exchange. Consider the example of a vote in a small committee, like a department meeting or a company board. Members typically have partial knowledge of other members' preferences, and before a decision is taken they will take strategic actions on which piece of information to disclose and to whom. There are important variables that they take into account: they know who are the people they can *count on*, they know who are the people whose opinion they *trust*, they also estimate whether their opinion will be *influential* and to whom. That is, they reason about the structure of various social networks that relate the individuals around them, and are able to devise and play complex strategies to achieve their goals and influence the result of the collective decision.

Similar phenomena are not restricted to collective decisions taken by small groups of people. Recent elections have shown how echo-chamber effects seem to polarise the opinions of societies, and how viral content can rapidly shift the public view, making traditional polls rather unreliable. Moreover, emerging technologies in the field of e-democracy present new challenges for designing trustworthy mechanisms for collective decisions on social media or on the Internet in general.

Social network analysis (Jackson, 2008; Easley and Kleinberg, 2010) is a burgeoning area which provides tools to analyse networks. It has proven very successful in many diverse fields, e.g., in the study of virus diffusion in biology, job market analyses, and targeted marketing. Well-established economic frameworks such as game theory the market analysis have started considering more complex models of society based on the study of networks. Social choice theory has been left relatively untouched by these new developments, and it is only recently that researchers have started to study the interplay between the networks relating members of a community and the collective decisions these individuals take.

This chapter provides an overview of novel approaches put forward in the computational social choice literature on the topic of social choice and social networks. References from outside of computer science are reduced to a minimum, but the problem is clearly of interest to political scientists and sociologists.

After briefly introducing some basic terminology and definitions in Section 9.2, we start in Section 9.3 by analysing the effects of social networks on standard collective decision scenarios. We first focus on the maximum-likelihood approach to voting, and then move to the setting of iterated strategic voting. Section 9.4 surveys a number of papers in which collective decision processes are designed to take into consideration the information coming from a network structure relating the individuals, in order to prevent strategic actions, or simply to obtain a better collective decision. In Section 9.5 we show how ideas from social choice theory have been used to devise novel models of opinion diffusion on social networks, with individual opinion updates defined by means of aggregation procedures. Section 9.6 concludes the survey pointing at a number of directions for future research.

9.2 Basic Definitions

We denote with $N = \{1, \dots, n\}$ the set of voters and $A = \{a, b, c, \dots\}$ the set of alternatives. We assume that voters in N are connected by a social network, represented as a directed graph identified by the set of its edges $E \subseteq N \times N$. Undirected graphs can be represented as symmetric directed graphs. We denote the neighbourhood of an individual i as $N(i) = \{j \in N \mid (j, i) \in E\}$, using a definition that can be applied to both directed and undirected networks. The structure of the network E can be constrained, e.g., to be a directed acyclic graph, a chain, a tree, or a more complex hierarchical structure. Influence or trust networks (Section 9.4.1) and delegation networks (Section 9.5) are assumed to be directed, while general social networks representing social acquaintances or information channels are assumed to be undirected. In the literature on social network analysis, graphs are typically undirected, and statistical assumptions are made on the distribution of the edges, for instance when considering random graphs or scale-free networks. While some of the papers referenced in this chapter are based on simulations and experiments, and hence make use of similar statistical assumptions, these properties are typically substituted by discrete graph properties in most algorithmic and theoretical papers.

Depending on the situation at hand, voters express their preferences in different forms. The classical approach from voting theory is that of a profile $P = (\succ_1, \dots, \succ_n)$, where $\succ_i$ is an irreflexive, complete and transitive binary relation over alternatives in A (i.e., a linear order). Another common setting is that of binary voting, where we will assume $A = \{0, 1\}$. A profile in this case is any $P = (p_1, \dots, p_n)$ where $p_i \in A$ is an individual ballot selecting one of the two alternatives. Further settings represent indidual opinions as real numbers, as vectors of binary views, or as propositional knowledge bases, and will be briefly introduced in the relevant sections.

A social choice mechanism will identify, in general, a procedure to obtain a collective decision from a profile of individual opinions. Examples can range from quota rules in binary voting, selecting 1 as the collective decision if the number of individual ballots for 1 exceeds a certain quota, or a rating system proposing a collective ranking of the alternatives from the evaluations of the individuals.

We recall here that there are two complementary views when assessing a social choice mechanism. In the first approach individuals are supposed to express their tastes over the set of alternatives, without there being any objective or true judgment on which alternative is best for the group. Political elections are a classic example, as well as lower stake decisions such as deciding which movie to watch among friends. Under this interpretation, the objective of the social choice mechanism is to guarantee the representativity of the collective choice with respect to the profile of individuals' ballots. The evaluation of a procedure under this interpretation is often done by means of axiomatic properties, with the most well-know result being Arrow's Theorem (Arrow, 1963).

The second approach views social choice as a problem of reconstructing an underlying correct ordering of the alternatives from noisy estimates received from the voters. These situations can occur, for instance, when the board of a company needs to judge on what project to invest in, or a committee on which candidate to choose for a given job, or crowdsourcing applications aiming at collectively classifying images or texts. Under this interpretation, a social choice mechanism is assessed by (a) devising a suitable noise distribution, and (b) proving that the mechanism maximises the probability of recovering the correct ordering of alternatives under the given noise distribution. The simplest and most well-known result in this research area is Condorcet's Jury Theorem (Condorcet, 1785): in binary voting, the majority rule is the maximum likelihood estimator for the noise model that assumes each voter to be correct with probability $p > 0.5$.

9.3 Effects of Social Networks on Collective Choices

A social network in a voting context acts as an information filter transforming a global view (the potential winner of an election, the distribution of preferences among voters, ...) into a multiplicity of local realities, each observed by a voter in the neighbourhood defined by the network. This first observation is well explained by a simple example known as the majority illusion (Lerman et al., 2016):

Example 9.1. *The citizens of a town have to vote in a two-candidate election: those supporting the first candidate are represented with full nodes in Figure 9.1, and those supporting the second candidate with empty nodes. A couple of days before the election takes place, a polling firm asks each voter which candidate she or he thinks will win. The election will be decided by majority, and each voter's reply to the poll is based on her private observation of other voters' opinions in her own neighbourhood. That is, we assume that voters are connected by a social network, and that they only have access to the opinions of their direct neighbours.*

The situation depicted in Figure 9.1 results in a surprising failure of the poll: while only 3 individuals out of 14 support the first candidate (the full nodes), a majority of the voters respond to the poll that it is the second candidate who will win the election. To see this, take for instance the rightmost voter, marked with a $*$: she can see three full nodes and one empty node, reporting a probable victory of

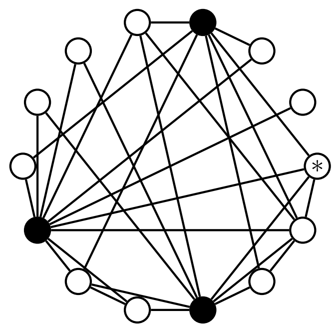

Figure 9.1: A network showing the "majority illusion".

the first candidate. The same happens to all 11 empty nodes in the graph, resulting in a poll which forecasts a victory of the first candidate with 11 votes against 3.

In this section we review a selection of recent papers from the literature on computational social choice that analyse how the information structure induced by a social network can affect the functioning of various voting mechanisms.

9.3.1 Noisy Votes

One of the first papers analysing the effects of social networks in computational social choice is the work of Conitzer (2012). This paper views voting as reconstructing a ground truth from noisy votes (i.e., the maximum likelihood estimator viewpoint described in Section 9.2), and reaches the following conclusion: if we see the probability that an individual estimates the correct alternative as an independent factor from the probability of being influenced by her neighbours' opinions, then the best mechanism to recover the ground truth simply ignores the existence of a network. Let us see this result in its simplest form.

We are in the setting of binary voting, with $A = \{0, 1\}$ as the set of alternatives, and we assume that one of the two alternatives in A is the correct one (i.e., the ground truth), and we denote it with $c \in A$. A population of voters N connected by an undirected network E receives a noisy distribution of opinions supporting either candidate 0 or candidate 1, generating a profile of binary opinions P. In the maximum likelihood approach we are interested in the probability of observing profile of votes P given that the correct alternative is c, i.e., estimating $Prob(P \mid c)$.

The first important assumption made by Conitzer (2012) is that such probability can be factored as $Prob(P \mid c) = \prod_{i \in N} f_i(p_i, P_{N(i)} \mid c)$, where p_i is agent i's opinion and $P_{N(i)}$ the restriction of profile P to i's influencers. That is, the overall probability of observing profile P can be factored into independent probability functions f_i, one for every individual, calculating the probability of observing opinion p_i and neighbouring opinions $P_{N(i)}$ given that the correct alternative is c. Assume now that the probability of observing an agent's ballot can be factored into (a) the probability of observing her ballot given that the correct alternative is c, and (b) the probability of observing her ballot given the profile of her influencers' opinions. In formulas, for every voter i there exist functions g_i and h_i such that $f_i(p_i, P_{N(i)} \,|\, c) = g_i(p_i \,|\, c) \cdot h_i(p_i, P_{N(i)})$.

We can now proceed to look for the maximum likelihood alternative, i.e., the $\hat{c} \in A$ maximising $Prob(P \,|\, \hat{c})$. By the formulation above, this is equivalent to maximising the product of functions h_i and g_i in c. However, since function h_i does not depend on the alternative c under consideration it can be ignored in the maximisation process. Hence, the maximum likelihood alternative can be found by only looking at functions g_i, which are independent from the structure of the network since they only depend on the current opinion p_i of individual i.

The assumptions that led to the previous result are clearly strong ones, in particular the one assuming that individuals form their views independently albeit being on a network. In a follow-up paper, Conitzer (2013) introduces a noise model for profile generation which has similarities with the models of opinion diffusion that will be introduced in Section 9.5. Such noise model is also defined for two alternatives, and assumes that the network E is such that each voter has an odd number of neighbours.[1] In the *independent conversation model*, voters have conversations with all of their neighbours, each one resulting in an argument in favour of one of the two alternatives, and will vote according to the majority of the arguments received.

In formulas, we associate with each edge $e = (i, j) \in E$ a random alternative $A_e \in \{0, 1\}$, representing the result of the conversation between i and j: if $A_e = 0$, for instance, the two voters after the conversation will end up with an argument in favour of 0. A profile of conversations $A_E = (A_e \mid e \in E)$ is the profile of alternatives $A_e \in A$ associated with edges $e \in E$, and $n(c, A_e) = |\{e \in E \mid A_e = c\}|$ is the number of edges associated with alternative $c \in A$. With every profile of conversations A_E, we can associate a profile of votes P such that $p_i = maj(\{A_{(j,i)} \mid j \in N(i)\})$, with each voter supporting the alternative for which they obtained a majority of arguments supporting it. Assuming that alternative A_e is the correct one with probability $p > 0.5$, and is the incorrect one with probability $1 - p$, we can obtain the probability of observing a profile P given that the correct alternative is c as $Prob(P \,|\, c) = \sum_{A_E} p^{n(c,A_E)} \times (1-p)^{|E|-n(c,A_E)}$, where A_E ranges over all profiles of conversations that are consistent with the observed profile, i.e., such that $p_i = maj(\{A_{(j,i)} \mid j \in N(i)\})$. The maximum likelihood alternative c is therefore the one that maximises the expression above. Considering all profiles A_E that are consistent with the observed profile of votes P clearly leads to an exponential explosion and to computationally intractable problems. For instance, computing the figure above is $\#P$-hard (Conitzer, 2013).

The independent conversation model explained above has been refined and extended in related work. Tsang et al. (2015) include the assumption that agents are more easily convinced by arguments supporting the correct alternative than by those supporting the incorrect one. Procaccia et al. (2015) include multiple alternatives and a more sophisticated process of individual response to conversations, testing whether two families of voting rules are *accurate in the limit*, i.e., whether their probability of recovering the ground truth tends to 1 when the

[1] In social choice theory it is commonplace to assume that a collective decision is taken by an odd number of individuals. This assumption is less natural when applied to a social network, but can still be considered as realistic when the number of individuals on the network is sufficiently large.

number of voters tends to infinity (an approach in line with the Condorcet's Jury Theorem mentioned in Section 9.2).

9.3.2 Iterative Voting

Strategic voting, i.e., the possibility of misrepresenting one's vote to influence the result of an election in one's favour, is typically considered as a one-shot strategic game with perfect information. Iterative voting relaxes some of these assumptions (see, e.g., Chapter 4 of this book). In this model, an initial profile of votes is gathered, let it be P^0. Such profile can be assumed to consist of the agents' truthful preferences, but can also be a completely random preference profile. Summary information about P^0 is then broadcast to voters, for instance in the form of a poll giving the percentage of voters supporting each of the candidates. Voters can then respond to this information by submitting a new vote, to form a new profile P^1. The process is repeated until a stable state is reached, i.e., a state in which no voter has an incentive to change her voting ballot. A social network in the setting of iterative voting has the effect of filtering the information available to voters, who will then take into consideration only the partial information observed in their neighbourhood when updating their vote.

Tsang and Larson (2016) generalise the model of iterative voting by giving each voter access only to the ballots of her neighbouring voters. By associating numerical values (cardinal preferences) to voters' ordinal preferences, they monitor experimentally the evolution of a number of parameters such as the price of stability and the price of honesty, as well the effects of iterative strategic voting on the ranking of less popular alternatives.

The same underlying principle is used by Sina et al. (2015): voters in the network can see how their neighbours voted at each stage, and this information can also be complemented by a summary poll obtained from the entire profile. The authors show that a central authority that is in control of the network structure is able to easily influence the result of an election. Technically, they present a polynomial-time algorithm that makes a chosen candidate the winner of the election by only adding a linear number of edges to the network. The existence of such a central authority should not be viewed as unrealistic: consider for instance the Facebook feed, that has the ability to block the information flow from certain connections or to suggest new ones.

9.3.3 Coalitional Games and Voting Equilibria

The presence of a social network has effects on a number of other collective decision processes. One example is studied in the work of Elkind (2014) and Igarashi (2017), who analyse coalitional games on a network in which coalition formation is restricted to connected subsets of players. A similar approach is taken by Igarashi et al. (2017), who study problems of group activity selection where the allocation of activities can only consider connected subgroups of individuals, as well as by Gourvès et al. (2016), who refine the notion of equilibrium in strategic voting by restricting deviations to coalitions of players that form a clique of the network and by including a form of empathy in the behaviour of the agents.

9.4 Social Choice Mechanisms over Networks

A social choice problem typically consists in taking a collective decision starting solely from the voters' preferences. When this input is complemented with knowledge of the social network that relates the voters, novel mechanisms can be conceived to take this additional information into consideration. In this section we survey a variety of procedures for collective choice that are designed to be implemented on a network of voters.

9.4.1 Liquid Democracy

The road from representative democracy, in which voters elect representatives that later take decisions on their behalf, to direct democracy, where voters directly take decisions on the issues at stake, is full of hybrid voting systems. One of these is *proxy voting* or *liquid democracy*.

Consider a set N of voters who need to take a decision on one single issue, and who are connected by an undirected social network E. Voters are allowed to either vote directly in favour (1) or against (0) the issue, or to *delegate* their vote to any other voter in their neighbourhood $N(i)$ who will act as their proxy.[2] The crucial assumption here is that delegations are transitive, i.e., a voter who received 4 delegations can in turn delegate her vote together with the extra 4 delegations to another voter.[3] The input of this particular social choice problem consists then of a profile P of ternary ballots, with each individual i specifying a vote for 0 or 1, or a proxy $j \in N(i)$, inducing a *delegation network* on N which we shall call E_P. Observe that every node in the delegation graph E_P has outdegree at most one, since every voter can delegate to only one person, restricting considerably the type of graphs that can be encountered in this setting.

There are several possibilities to take a collective decision on such input profiles. If the principle of "one man, one vote" needs to be kept, then a weighted majority rule will be used, defining the weight of a proxy as the number of voters that have delegated their vote, either directly or via a chain of delegations, to the proxy. Formally, let a proxy be a voter i that expressed a direct vote either for 0 or for 1. Its weight can then be computed as $w_i = |\{j \in N \mid j\ E_P^*\ i\}|$, where jE_P^*i is the reflexive and transitive closure of the delegation graph E_P. Christoff and Grossi (2017b) analyse this setting through the lens of judgment aggregation, assuming that there are multiple interconnected binary issues at stake. Among the numerous problems they consider, there is the existence of cycles of delegations, which may cause a failure of the "one man, one vote" principle. They also observe that computing the result of a liquid democracy vote can be performed by finding a stable state of a suitable opinion diffusion model on the delegation network, a problem that we will analyse in detail in Section 9.5.

An alternative direction is to use off-the-shelf spectral ranking techniques to

[2]The assumption that delegates must be part of the social neighbourhood of a voter can be relaxed, but it is often assumed in practice to guarantee a minimal level of trust in the delegation process.

[3]Non-delegable proxy voting has been studied in the social choice literature by, e.g., Miller (1969). Moreover, recent work in computational social choice assess the use of non-delegable proxies in elections with a small number of active voters (Cohensius et al., 2017).

compute the weight of each proxy on the delegation network. This includes the well-known PageRank algorithm (Boldi et al., 2009) and the Katz index (Boldi et al., 2011). In the latter paper, the authors introduce a damping factor to limit the effect of long chains of transitive delegations. Their idea is the following: every delegation path leading to a proxy contributes to its weight with a sum that is inversely proportional to the path length. Clearly, the "one man, one vote" principle is lost, but in various settings (such as voting on online social media) the trust relation underlying a delegation may not be strong enough to justify the transfer of voting power. If $\alpha \in (0,1)$ is a damping factor, and we denote with $|p|$ the length of a path p,[4] we can define the weight of a proxy i as $w_i = \sum_{p \in \mathit{Path}(-,i)} \alpha^{|p|}$, where *Path*$(-,i)$ is the set of delegation paths ending in i. Once the weights of the proxies have been computed, a vote by weighted majority can be staged to compute the winning alternative. These methods have also been proposed and tested to construct a personalised recommender system (Boldi et al., 2015).

Non-binary versions of proxy voting give rise to a variety of other problems, as pointed out by the work of Behrens et al. (2014). For instance, when a large number of alternative proposals need to be considered and voted on, the ordering in which these competing alternatives are presented to the voters is of crucial importance. Skowron et al. (2017) mention this problem as a potential application of their work on proportionality, in which a ranking of alternatives needs to be constructed from individual approval ballots.

9.4.2 Ratings and Recommendations

Obtaining a collective rating for products or objects (classical example: the rating application TripAdvisor) is a collective decision problem that is close to that of obtaining recommendations for users (classical example: Netflix or Amazon's recommender systems). Both problems are typically solved at a global scale, considering preferences and calculating similarities over the whole of the users' data. In this section we look at three examples of how a social network can play a role in obtaining more meaningful and robust recommendations.

Let us first set some definitions. As before, let N be a set of individuals connected by a social network E. A subset of individuals $V \subseteq N$, called voters, express their opinions over a set of alternatives A, which we assume by simplicity consists of a single object. These opinions, which we denote $opinion(i)$ for $i \in N$, can take different forms, e.g., a like/dislike, a numerical rating, or an evaluation on a discrete scale. The problem of recommendation is then the following: given a target non-voter $v^* \in N \setminus V$ and the vector of individual opinions from voters in V, should the product be recommended to v^*?

One of the simplest solutions is to first calculate a collective rating for the object, perhaps as the average or the median of all the ratings expressed by voters, and recommend the object to the non-voter if the collective rating exceeds a certain threshold. This approach is not ideal for a number of reasons, one of which being its vulnerability to a variety of attacks by strategic agents. Consider for instance the problem of *false-name manipulations*: by opening multiple

[4]For ease of explanation we discard the case of infinite paths.

fake accounts providing the same rating as her own, a voter is able to influence both the median and the average towards her evaluation. When the social network connecting the individuals is known, a simple idea can however be used to obtain more robust rating mechanisms: instead of computing a global collective rating, recommendations should be based on a personalised form of rating, computed from the opinions of those voters that belong to a suitable notion of neighbourhood of the target user v^*. We now review three papers that formalised and studied a personalised approach to ratings and recommendations.

Andersen et al. (2008) carry out an axiomatic study of trust-based recommender systems, assuming that voters express binary opinions (like/dislike) on a single object. Let v^* be a non-voter and v a voter, and let $Prob(v^*, v)$ be the probability of reaching voter v from v^* with a random walk on the trust network E. Their *random walk* system will then recommend the product to v^* if the probability of reaching a voter from v^* whose opinion is "like" is higher than the probability of reaching a voter whose opinion is "dislike". Hence, they base recommendations on the ratings of all individuals in the connected component of E containing v^*, weighted by the network structure.

Grandi and Turrini (2016) focus instead on real-valued or discrete-scale opinions on a single object. Rather than presenting the overall average rating to all agents, the authors propose the use of *personalised ratings*, obtained by computing for each agent i the average rating of her direct neighbours on the trust network E. The strategic problem studied is that of bribery, in which an external agent tries to influence the rating of the object by bribing individuals to increase their ratings. Compared to providing the overall average of ratings, personalised ratings make bribery by an external agent more costly and, under some specific assumptions, not profitable.

Brill et al. (2016a) refine this idea by considering a broader notion of neighbourhood than the direct connections on the trust network. Their idea is the following one: let $F(u)$ be the set of users that are disconnected from the target user v^* as a result of removing node u from the network. A voter v is called *legitimate* if it does not belong to $F(u)$ for any $u \in N$, that is, if v is still reachable from v^* after the deletion of any node. They propose a recommender system based on this notion of neighbourhood, showing that it is false-name strategy-proof, i.e., it cannot be manipulated by the addition of fake users.

9.4.3 Social Polls and Empathetic Preferences

Gaspers et al. (2013) study the computational complexity of determining possible and necessary winners in the context of social polling. The key aspect of their model is that agents vote following a sequential order, observing the ballots that were previously cast by voters in their direct neighbourhood.

While voters are typically assumed to be influenced by the decisions of their neighbours, Salehi-Abari and Boutilier (2014) propose a different model in which voters' utilities on collectively decided alternatives depend positively on the utilities of their neighbours, to represent voters' desire to see others satisfied with the chosen alternative.

9.5 Opinion Diffusion

As in the introductory example of a committee decision, the addition of a network in a social choice problem is typically done to take into consideration phenomena such as social influence and the information diffusion that takes place before (or during) the decision process. The literature on social network analysis abounds with models of influence-based opinion diffusion. Two classical examples are threshold models (Granovetter, 1978), with more recent generalisations by Kempe et al. (2003, 2005), and the De Groot or Lehrer-Wagner model (de Groot, 1974; Lehrer and Wagner, 1981). With the notable exception of the recent work of Friedkin et al. (2016), these models are based on a simple representation of individual opinions as either a binary view on a single issue, or a real-valued opinion in the interval $[0, 1]$. We do not survey this literature here, pointing at the classical references in social network analysis for a detailed survey (Jackson, 2008; Easley and Kleinberg, 2010).

In this section, instead, we focus on a number of recent papers that borrowed techniques from social choice theory and knowledge representation to model the diffusion of *complex* opinions: multiple binary issues, linear orders, and belief bases. Such models are founded on the observation that influence on a social network is itself a social choice problem, since every individual uses some form of aggregation when updating her opinion based on the opinions of her influencers (e.g., her direct neighbours). This gives rise to a discrete-time iterative process in which at each point in time a number of agents on the network update their opinion using an aggregation procedure F_i that takes into account the opinions of the neighbours as well as her own one. Typical problems that can be studied are the *termination* of the iterative processes, characterising those networks or initial profiles of opinions on which the process is guaranteed to reach a stable state, and *convergence*, obtaining a characterisation of termination states. A typical example of the latter problem is convergence to consensus, i.e., the identification of properties of the network that guarantee that all agents will have unanimous opinions at the end of the diffusion process.

Example 9.2. *Consider the following example from Brill et al. (2016b). Let there be four agents on the influence network described below, and let each agent express her preferences in the form of a linear order over three alternatives* a, b *and* c*:*

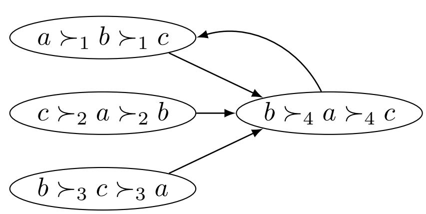

Figure 9.2: The influence of a Condorcet cycle.

The preferences of agents $1, 2,$ *and* 3 *form what is known as a* Condorcet cycle*, i.e., the majority relation of their preferences is cyclic. Assume now that opinion update process follows the opinion of the majority of an agent's influencers, by swapping* *adjacent pairs of alternatives in her preferences accordingly. In the example above, we can devise a sequence of* asynchronous *(i.e., sequential) updates that terminates (albeit not to consensus). First, we let agent* 4 *update on pair* ab*, moving to preference* $a \succ_4 b \succ_4 c$*. After that, no further updates are possible: even though agent* 4 *disagrees with its influencers on pair* ac*, this pair cannot be swapped since it is no longer adjacent in* $\succ_4$*. In the same example, if we consider* synchronous *updates by the agents, then we can devise a sequence that does not terminate, namely the one where agents* 1 *and* 4 *update repeatedly on pair* ab*. Observe that in any update sequence agents* 2 *and* 3 *never update their preferences.*

Let us start from the case of individual opinions over *multiple binary issues*. Consider a set of agents N on a directed network E, each one having a binary opinion over a set of issues I. Let for instance agent i express her opinion as $p_i = (0, 1, 0, 0)$ on a domain composed of four issues. At each point in time, an individual i has access to the opinions of its direct neighbours or influencers in $N(i)$, and updates her opinion using an aggregation function F. If P^t is the profile of individual profiles at time t, then $p_i^{t+1} = F(P^t_{|N(i)}, p_i^t)$, where $P^t_{|N(i)}$ is the restriction of profile P to the individuals in $N(i)$. The process can be *synchronous*, when all individuals update at the same time, or *asynchronous*, when individuals update one after the other. Grandi et al. (2015) provide algorithms and characterisation results for the termination of synchronous opinion diffusion models on multiple binary issues described above. Some of these results have been subsumed by the recent work of Christoff and Grossi (2017a), which focuses on aggregation procedures F that satisfy some natural choice-theoretic properties. When the aggregation procedure F is a quota rule (Dietrich and List, 2007), these models are equivalent to Granovetter's threshold functions. Goles and Olivos (1980) show that under such assumptions, every sequence of synchronous updates always terminates to a stable state or cycles with period 2. Different opinion diffusion models over multiple binary issues can be defined from any rule from the judgment aggregation literature (Endriss, 2016). Slavkovik and Jamroga (2016), for instance, explore the use of a distance-based procedure to reach consensus on a network. Finally, the addition of an integrity constraint relating the multiple issues is not a trivial generalisation, and initial results in this directions have been obtained by Botan et al. (2017).

In a voting setting, individual opinions are typically assumed to be *linear orders* over a set of alternatives, inducing more complex diffusion models studied by Brill et al. (2016b) and by Farnoud et al. (2013) (the latter for the case of a complete network). The main problem faced in these models is the presence of intransitive majorities, as presented in Example 9.2: three influencers with preferences $a \succ_1 b \succ_1 c$, $c \succ_2 a \succ_2 b$, and $b \succ_3 c \succ_3 a$, influencing a fourth individual with preferences $b \succ_4 a \succ_4 c$. Aggregating the three linear orders by majority results in a cycle: how should then the fourth agent update her preferences? One possible solution is to restrict individual updates to swapping pairs that are already adjacent in the individual ordering that is being updated. For instance, in the previous example, the fourth agent may update on pair (b, a), switching to $a \succ_4 b \succ_4 c$, which agrees on all adjacent pairs with the (intransitive) majority of her influencers. In this model, Brill et al. (2016b) provide a termination result on arbitrary networks under additional assumptions on the profile of individual preferences, and show that on directed acyclic graphs the diffusion model preserves classical domain restrictions from voting theory such as single-peakedness.

Belief bases are sets of formulas in propositional logic that are used to compactly represent the beliefs of an agent about the current state of the world. For example, if p stands for "it is raining" and q stands for "it is cold", an agent with a belief base of $p \vee q$ believes that it is raining or it is cold, or both. These mathematical objects have been used as individual opinions by Schwind et al. (2015) and Cholvy (2016), to define diffusion processes in line with those described above. Individuals on a network are assumed to have access to the belief bases of their direct neighbours, and to use this information to update their current belief base using a belief merging operator (see, e.g., Chapter 7 of this book). In particular, Schwind et al. (2015) analyses the axiomatic properties of this model, varying the belief merging technique that is used for the update of individuals' belief bases.

The strategic aspects of the diffusion models described above are of clear interest, both for individual strategic actions such as misrepresenting one's own opinion, and for external actions such as bribery or control. Some of these problems are just beginning to be explored in the case of belief bases (Schwind et al., 2016) and binary issues (Grandi et al., 2017; Bredereck and Elkind, 2017).

9.6 Conclusions

In this chapter we surveyed recent work in the computational social choice literature on social choice and social networks.

First, we saw how social-network-related phenomena, such as social influence and an asymmetric distribution of information, can impact the result of standard procedures for collective decision-making. When taking a maximum likelihood approach to social choice, the structure of a network can be exploited to create novel noise models and new maximum likelihood estimators. An open question is then whether opinion diffusion models such as those defined in Section 9.5 can be interpreted as noise models, and what are the maximum likelihood estimators for them. The setting of iterative voting, when voters respond iteratively to a sequence of polls or elections, is also affected by individuals responding to local information filtered through the network. Political elections provide numerous examples to observe the consequences of network-related processes, from the majority illusion discussed in Section 9.3, to echo-chamber effects, to polarisation. Assessing the effects of social networks on the ability of taking collective decisions in society is a crucial topic for modern social choice, and a rich source of computational problems.

We then considered the problem of designing mechanisms for collective choice that are implemented on networks of voters. The classical example here is proxy voting, in which voters can delegate their voting power to a neighbour, inducing a delegation network. Various notions from spectral ranking can be used to compute the weights of voters and arrive at a collective decision. While classical weight functions from social network analysis have been tested, novel measures may be defined that are specific to a voting context. We also saw applications to the problem of obtaining ratings and recommendations of objects to users. In this setting, suitable notions of neighbourhood on the network can be used to provide personalised ratings and recommendations that are more robust against malicious strategic actions by both users and external agents.

Finally, we showed how social choice can contribute to the definition of novel models of opinion diffusion, based on the idea that individuals aggregate the views they receive from their neighbours on a social network. Depending on the application at hand, these models can be constructed using voting rules from classical social choice theory, aggregation procedures from judgment aggregation, or belief merging operators. These models present a number of open algorithmic challenges, most notably the characterisation of networks guaranteeing termination, and provide a computation-friendly representation of diffusion whose effects on social choice methods still need to be assessed.

Strategic aspects of collective decision-making on social networks are still largely unexplored. Agents may have multiple actions available, from adding or severing links on the network, to misrepresenting their opinion in different ways to different agents, to exercising their influence at various degrees. This new layer of strategic reasoning may have a significant impact on the problem of equilibrium selection in voting games, and we have seen some first studies in this direction in the area of iterative voting. Many of the ideas discussed in this chapter also have the potential to contribute to real-world applications: from ratings and recommendations on networks, to the rise of platforms for democracy and online decision making (see, e.g., Chapter 20 of this book).

Acknowledgments

The author gratefully acknowledges Markus Brill, Davide Grossi, and Arianna Novaro for their useful comments and discussions.

Bibliography

R. Andersen, C. Borgs, J. Chayes, U. Feige, A. Flaxman, A. Kalai, V. Mirrokni, and M. Tennenholtz. Trust-based recommendation systems: An axiomatic approach. In *Proceedings of the 17th International Conference on World Wide Web (WWW)*, 2008.

K. J. Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 2nd edition, 1963. First edition published in 1951.

J. Behrens, A. Kistner, A. Nitsche, and B. Swierczek. *Principles of Liquid Feedback*. Interaktive Demokratie e. V., 2014.

P. Boldi, F. Bonchi, C. Castillo, and S. Vigna. Voting in social networks. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management (CIKM)*, 2009.

P. Boldi, F. Bonchi, C. Castillo, and S. Vigna. Viscous democracy for social networks. *Communications of the ACM*, 54(6):129–137, 2011.

P. Boldi, C. Monti, M. Santini, and S. Vigna. Liquid FM: recommending music through viscous democracy. In *Proceedings of the 6th Italian Information Retrieval Workshop*, 2015.

S. Botan, U. Grandi, and L. 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.

R. Bredereck and E. Elkind. Manipulating opinion diffusion in social networks. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.

M. Brill, V. Conitzer, R. Freeman, and N. Shah. False-name-proof recommendations in social networks. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2016a.

M. Brill, E. Elkind, U. Endriss, and U. Grandi. Pairwise diffusion of preference rankings in social networks. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016b.

L. Cholvy. Diffusion of opinion and influence. In *Proceedings of the 10th International Conference on Scalable Uncertainty Management (SUM)*, 2016.

Z. Christoff and D. Grossi. Stability in binary opinion diffusion. In *Proceedings of the 6th International Conference on Logic, Rationality, and Interaction (LORI)*, 2017a.

Z. Christoff and D. Grossi. Binary voting with delegable proxy: An analysis of liquid democracy. In *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2017b.

G. Cohensius, S. Mannor, R. Meir, E. A. Meirom, and A. Orda. Proxy voting for better outcomes. In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2017.

Condorcet, Marie Jean Antoine Nicolas de Caritat, Marquis de. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris, 1785.

V. Conitzer. Should social network structure be taken into account in elections? *Mathematical Social Sciences*, 64(1):100–102, 2012.

V. Conitzer. The maximum likelihood approach to voting on social networks. In *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing*, 2013.

M. H. de Groot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

F. Dietrich and C. List. Judgment aggregation by quota rules: Majority voting generalized. *Journal of Theoretical Politics*, 19(4):391–424, 2007.

D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World.* Cambridge University Press, 2010.

E. Elkind. Coalitional games on sparse social networks. In *Proceedings of the 10th International Conference on Web and Internet Economics (WINE)*, 2014.

U. Endriss. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17. Cambridge University Press, 2016.

F. Farnoud, E. Yaakobi, B. Touri, O. Milenkovic, and J. Bruck. Building consensus via iterative voting. In *Proceedings of the 2013 IEEE International Symposium on Information Theory*, 2013.

N. E. Friedkin, A. V. Proskurnikov, R. Tempo, and S. Parsegov. Network science on belief system dynamics under logical constraints. *Science*, 354(6310):321–326, 2016.

S. Gaspers, V. Naroditskiy, N. Narodytska, and T. Walsh. Possible and necessary winner problem in social polls. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2013.

E. Goles and J. Olivos. Periodic behaviour of generalized threshold functions. *Discrete Mathematics*, 30(2):187–189, 1980.

L. Gourvès, J. Lesca, and A. Wilczynski. Strategic voting in a social context: Considerate equilibria. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, 2016.

U. Grandi and P. Turrini. A network-based rating system and its resistance to bribery. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

U. Grandi, E. Lorini, and L. Perrussel. Propositional opinion diffusion. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.

U. Grandi, E. Lorini, A. Novaro, and L. Perrussel. Strategic diffusion of opinions on a social network. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2017.

M. Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978.

A. Igarashi. Supermodular games on social networks. In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2017.

A. Igarashi, D. Peters, and E. Elkind. Group activity selection on social networks. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.

M. O. Jackson. *Social and Economic Networks*. Princeton University Press, 2008.

D. Kempe, J. M. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2003.

D. Kempe, J. M. Kleinberg, and E. Tardos. Influential nodes in a diffusion model for social networks. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP)*, 2005.

K. Lehrer and C. Wagner. *Rational Consensus in Science and Society*. Springer, 1981.

K. Lerman, X. Yan, and X.-Z. Wu. The "majority illusion" in social networks. *PLoS ONE*, 11(2):1–13, 02 2016.

J. C. Miller. A program for direct and proxy voting in the legislative process. *Public Choice*, 7(1):107–113, 1969.

A. D. Procaccia, N. Shah, and E. Sodomka. Ranked voting on social networks. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.

A. Salehi-Abari and C. Boutilier. Empathetic social choice on social networks. In *International conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2014.

N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Belief revision games. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Is promoting beliefs useful to make them accepted in networks of agents? In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

S. Sina, N. Hazon, A. Hassidim, and S. Kraus. Adapting the social network to affect elections. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.

P. Skowron, M. Lackner, M. Brill, D. Peters, and E. Elkind. Proportional rankings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.

M. Slavkovik and W. Jamroga. Iterative judgment aggregation. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, 2016.

A. Tsang and K. Larson. The echo chamber: Strategic voting and homophily in social networks. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2016.

A. Tsang, J. A. Doucette, and H. Hosseini. Voting with social influence: Using arguments to uncover ground truth (extended abstract). In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.