Social Network Analysis for Judgment Aggregation

Paper 487

ABSTRACT

Judgment aggregation investigates the problem of how to aggregate several individuals' judgments on some logically connected propositions into a consistent collective judgment. The majority of work in judgment aggregation is devoted to studying impossibility results, but the (social) dependencies that may exist between voters is traditionally not studied. In this paper, we use techniques from social network analysis to study the relations between the individuals participating in a judgment aggregation problem by analysing the similarity between their judgments in terms of social networks. We show that using these techniques we obtain a more fine-grained approach than the *prototype* voter rule in addition to identifying closely related groups of agents sharing some common interests concerning the given issues. We provide both theoretical and empirical results supporting our claims, in addition to analysing the complexity of the approach adopted.

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multi-agent systems, Cooperation and coordination*

General Terms

Experimentation, Theory

Keywords

Artificial social systems, Social and organizational structure, Collective decision making

1. INTRODUCTION

The interactions within a multiagent system include cooperation and coordination. To be able to coordinate and cooperate, intelligent agents need to reach collective consents, namely binding group decisions, over issues such as beliefs, action and desires. Recently, it has been recognized that such multiagent problems can benefit from social choice methods (see e.g. [24] for an overview). Social choice theory studies the problem of reaching collective consent between a group of people in the scope of economic theory. It includes voting theory, preference aggregation and judgment aggregation.

The problem of judgment aggregation investigates how to aggregate individual judgments on logically related propositions to a group judgment on those propositions. Examples of groups that need to aggregate individual judgments can be expert panels, legal courts, boards, and councils. The problem of aggregating judgments started to attract considerable attention in the last ten years, since it has been shown to be general in the sense that it subsumes voting theory and preference aggregation [16].

The majority of work in judgment aggregation is devoted to studying impossibility results in the style of the work in preference aggregation by Arrow [1, 17], leading to the development of several aggregation rules such as majority outcome, premise-based aggregation, and conclusion-based aggregation [16]. These rules are all concerned with the general problem of selecting outputs that a *consistent* or *compatible* [15] with individual judgments. However, the (social) dependencies that may exist between the voting agents and how this influences the outcome of the voting scenario is largely neglected.

A representation of the social structure of a judgment aggregation problem makes it possible to identify influential voters in the entire group or in a subgroup of voters. Such information can for instance be used to detect cartels in voting scenarios, or to simplify the voting problem (i.e. reduce the number of voters), to the extreme case where a single agent is chosen to represent the entire group. It seems that the well-known science fiction writer Isaac Asimov had something similar in mind when he was writing the story "Franchise", in which the election process in the United States is controlled by super-computer Multivac, who selects a single representative individual from the entire population as the representative voter [2]:

> *"The man presented credentials, stepped into the house, closed the door behind him and said ritualistically, 'Mr. Norman Muller, it is necessary for me to inform you on the behalf of the President of the United States that you have been chosen to represent the American electorate on Tuesday, November 4, 2008."*

It does not seem obvious to extract such information from a voting scenario, merely by relying on the votes of the indi-

Appears in: *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), Lomuscio, Scerri, Bazzan, Huhns (eds.), May, 5–9, 2014, Paris, France.*

vidual agents. However, we believe that a possible natural solution to this problem can be provided by using techniques from social network analysis (SNA) to characterise dependencies between voters. SNA views social relationships in terms of graph theory, consisting of nodes (representing individual agents within the network) and ties (which represent relationships between the agents, such as friendship, kinship, organisational position, sexual relationships, etc.) [22, 10].

In this paper, we explore the possibility to apply SNA to judgment aggregation by translating a judgment aggregation problem to two social networks, depicting respectively agreement between agents and correlation between issues. We analyse the first network using three common measures of node centrality from SNA: degree, closeness, and betweenness. We proof an equivalence between the prototype voting rule and an instance of the degree centrality measure. We further investigate the possibility to select representative voters from a judgment aggregation scenario using these measures by performing experiments and obtain promising results, providing strong evidence for the use of using tools from SNA in order to analyse judgment aggregation scenarios.

The paper is organised as follows: We start by discussing related work in Section 2. In Section 3 we introduce the basic notitions of judgment aggregation and several voting rules, and we introduce basic terminology from social network analysis in Section 4. In Section 5 we show how we can obtain two social networks from a judgment aggregation problem using matrix operations. We use these networks in Section 6, where we show correspondence between properties of the graph and different measures in judgment aggregation. Empirical results and a complexity analysis are discussed in Section 7; Future work and conclusions are given in Section 8.

2. RELATED WORK

There have been several approaches that have combined notions from graph theory with social choice theory. For instance, [12] considers the problem of graph aggregation, where individuals do not give a judgment over alternatives, but instead provide a directed graph over a common set of vertices. Judgment aggregation reduces then to computing a single graph that best represents the information inherent in this profile of individual graphs. This is considerably different from our work, since we obtain a graph from the dependencies between voters, assuming that voters give a judgment over alternatives. [5] proposes a qualitative graphical representation of preferences that reflects conditional dependence and independence of preference statements under a ceteris paribus (all else being equal) interpretation. This does not take into account the dependencies between individuals that may exist based on their judgment, but only those between preferences. Salehi-Abari and Boutilier [23] take the opposite direction than we do: Where we take social choice theory as the starting point and apply social network theory to it, they propose to apply social choice techniques to social networks where agents derive utility based on both their own intrinsic preferences and the satisfaction of their neighbors. In [9], each voter is assume to have an argumentation network that motivates its judgments, which can then be used to select outcomes which are not just consistent, but compatible with the individuals argumentation network. Our work does not assume any additional information from the voter, which makes it considerably different.

3. JUDGMENT AGGREGATION

In this section we recall the framework of judgment aggregation [16, 24], which we shall be working with. We will formulate the judgment aggregation problem as *binary aggregation with integrity constraints*, which is shown to be equivalent to judgment aggregation when the individual judgments of the agents are complete and consistent [14]. Besides introducing the basic framework, we also define several voting rules that we will use throughout the paper.

3.1 Basic Definitions

Let $\mathcal{N} = \{1, 2, \ldots, n\}$ be a finite set of agents, and let $\mathcal{I} = \{1, 2, \ldots, m\}$ be a finite set of *issues*. We want to model collective decision making problems where the group of agents $\mathcal{N}$ have to jointly decide for which issues in $\mathcal{I}$ to choose "yes" and for which to choose "no". A ballot $B \in \{0,1\}^m$ associates either 0 ("no") or 1 ("yes") with each issue in $\mathcal{I}$. We write B_j for the jth element of B. Thus, $B_j = 1$ denotes that the agent has accepted the jth issue, and $B_j = 0$ denotes that the agent has rejected it.

In general, not every possible ballot might be a *feasible* or *rational* choice. For instance, if the issues are tasks that are to be executed by a group of agents, then a task constraint might mean that deciding to execute certain tasks makes it impossible to execute other tasks.

Formally, let $PS = \{p_1, \ldots, p_m\}$ be a set of propositional symbols, one for each issue $\mathcal{I}$. An *integrity constraint* is a formula $IC \in \mathcal{L}_{PS}$, where $\mathcal{L}_{ps}$ is obtained from PS by closing under the standard propositional connectives. Let $Mod(IC) \subseteq \{0,1\}^m$ denote the set of models of IC, i.e. the set of rational ballots satisfying IC.

A *profile* is a vector of rational ballots $\mathbf{B} = (B_1, \ldots, B_n) \in Mod(IC)^n$, containing one ballot for each agent. We write B_{ij} to denote the ith agent's choice about the jth issue, i.e. the jth choice of ballot B_i. Since ballots are vectors themselves, we can consider $\mathbf{B}$ as a matrix of size $n \times m$. The *support* of a profile $\mathbf{B} = (B_1, \ldots, B_n)$ is the set of all ballots that occur at least once within $\mathbf{B}$:

$$\textsc{Supp}(\mathbf{B}) = \{B_1\} \cup \ldots \cup \{B_n\}.$$

A *voting rule* $F : \{0,1\}^{m \times n} \to 2^{\{0,1\}^m}$ is a function that maps each profile $\mathbf{B}$ to a set of ballots. This means that an aggregation rule can have one or multiple outcomes, also called an *irresolute voting rule*. A voting rule is called *collectively rational* when all outcomes satisfy the integrity constraints.

One of the most well-known voting rules is the (weak) *majority rule*, which accepts an issue if a weak majority of the agents accept it:

$$Maj(\mathbf{B})_j = 1 \text{ iff } |\{i \in \mathcal{N} \mid \mathbf{B}_{ij} = 1\}| \geq \left\lceil \frac{n}{2} \right\rceil.$$

Example 3.1. *Suppose a judgment aggregation scenario consisting of six agents (a, b, c, d, e, f) voting on an agenda composed of four issues (p, q, r, z). The agenda is subject to the following integrity constraint: $IC = (p \wedge q \wedge r) \Leftrightarrow z$.*

Issue:	p	q	r	z
a	0	1	1	0
b	1	0	0	0
c	1	1	1	1
d	1	0	0	0
e	1	0	1	0
f	0	0	1	0
Maj	1	0	1	0

3.2 The Distance-Based Rule

The Hamming distance between two ballots $B = (B_1, \ldots, B_m)$ and $B' = (B'_1, \ldots, B'_m)$ is defined as the sum of the amount of issues on which they differ:

$$H(B, B') = |\{j \in \mathcal{I} \mid B_j \neq B'_j\}|$$

For example, $H((1,0,0),(1,1,1)) = 2$. The Hamming distance between a ballot B and a profile $\mathbf{B}$ is the sum of the Hamming distances between B and the ballots in $\mathbf{B}$:

$$\mathcal{H}(B, \mathbf{B}) = \sum_{i \in \mathcal{N}} H(B, B_i)$$

Definition 1. *Given an integrity constraint IC, the distance-based rule DBR is the following function:*

$$DBR^{IC}(\boldsymbol{B}) = \underset{B \in Mod(IC)}{\text{argmin}} \sum_{i \in \mathcal{N}} H(B, B_i)$$

Note that the DBR^{IC} is collectively rational by definition because it only considers outcomes in $Mod(IC)$. It has good social choice-theoretic properties and is one of the most studied rules in preference aggregation. However, it has a rather high complexity of Θ_2^p.

3.3 The Prototype Voter Rule

Grandi and Pigozzi propose in [15] a way to minimize the complexity of computing an outcome of a judgment aggregation problem. The space of the possible outcomes is reduced by taking into consideration as possibilities only the ballots proposed by the voters. This way of selecting an outcome is a generalised dictatorship and proposed in judgment aggregation under the name *prototype rule* [18]:

Definition 2 (Prototype Rule). *The prototype rule is the voting rule that selects those individual ballots that minimise the Hamming distance to the profile:*

$$PRO(\boldsymbol{B}) = \underset{B \in \text{Supp}(B)}{\text{argmin}} \ \mathcal{H}(B, \boldsymbol{B})$$

The complexity of winner determination for PRO is in $O(mn \log n)$, which is considerably better than the DBR, but the outcomes are less optimal they are restricted to the ballots of the voters.

4. SOCIAL NETWORK ANALYSIS

A social network usually is represented as a graph. The vertices are the individuals, and the edges represent the social connections. In this paper, we consider the symmetric case where social networks are represented by undirected graphs. An edge which joins a vertex to itself is called a *loop*. The number of edges that are incident to a vertex is called the *degree* of a vertex. The *neighborhood* of a vertex v is the set of all vertices adjacent to v.

We denote a weighted network (or weighted graph) with $G = (V, E, W)$ with the vertex set $V(G) = \{v_1, \ldots, v_n\}$, edge set E, and weight matrix W, where each edge $e = (v_i, v_j)$ is labeled with a weight w_{ij}. We assume that if two vertices are not connected, then there exists an edge of weight 0 connecting them. Since we only consider undirected networks, $w_{ij} = w_{ji}$. We define the *sum-weight* s_i of a vertex v_i with $s_i = \sum_{j=1}^{n} w_{ij} = \sum_{u \in N(v_i)} w_{v_i u}$, where $N(v_i)$ is the neighborhood of v_i. We denote the degree k_i of a vertex v_i with $k_i = |N(v_i)|$, i.e. k_i denotes the number of neighbors of v_i.

The centrality of vertices, identifying which vertices are more ”central” than others, has been a key issue in network analysis. Freeman [13] originally formalized three different measures of vertex centrality: degree, closeness, and betweenness. Degree is the number of vertices that a focal vertex is connected to, and measures the local involvement of the vertex in the network. To also take the global structure of the network into account, closeness centrality was defined as the inverse sum of shortest distances to all other vertices from a focal vertex. A main limitation of closeness is the lack of applicability to networks with disconnected components: two vertices that belong to different components do not have a finite distance between them. Betweenness provides a solution by assessing the degree to which a vertex lies on the shortest path between any two other vertices.

These measures are originally formalised for binary graphs [13], but we will consider recent proposal [21] that uses a tuning parameter α to control the relative importance of number of edges compared to the weights on the edges. We now introduce the three centrality measures for degree (C_D), closeness (C_C), and betweenness (C_B).

The degree centrality measure is defined as the product of the number of vertices that a focal vertex is connected to, and the average weight to these vertices adjusted by the tuning parameter. The degree centrality for a vertex i is computed as follows:

$$C_D^{W\alpha}(i) = k_i \times \left(\frac{s_i}{k_i}\right)^{\alpha} = k_i^{(1-\alpha)} \times s_i^{\alpha} \tag{1}$$

where W is the weight matrix of graph, α is a positive tuning parameter, k_i is the size of the neighborhood of vertex i and s_i the sum of the weights of the incident edges. If α is between 0 and 1, then having a high degree is favorable over weights, whereas if it is set above 1, a low degree is favorable over weights. In Section 6 we elaborate on different levels of α for degree centrality.

The closeness and betweenness centrality measures rely on the identification and length of the shortest paths among vertices in the network. Dijkstra [11] proposed an algorithm that finds the path of least resistance, and was defined for networks have the weights represented costs. Since weights in our network are operationalisations of tie strength and not the cost of them, the tie weights need to be reversed before directly applying Dijkstra's algorithm to identify the shortest paths in these networks:

$$d^{W\alpha}(i,j) = \min\left(\frac{1}{(w_{ih})^{\alpha}} + \ldots + \frac{1}{(w_{hj})^{\alpha}}\right) \tag{2}$$

where h are intermediary vertices on paths between vertex i and j and α is again a positive tuning parameter. Again, α controls the relative important of edge weight compared to number of edges. We will return to the analysis of α for closeness and betweennees in Section 7.

We define closeness centrality and betweenness centrality as follows:

$$C_C^{W\alpha}(i) = \left[\sum_{j}^{N} d^{W\alpha}(i,j)\right]^{-1} \quad (3)$$

$$C_B^{W\alpha}(i) = \frac{g_{jk}^{W\alpha}(i)}{g_{jk}^{W\alpha}} \quad (4)$$

where g_{jk} is the number of binary shortest paths between two vertices, and $g_{jk}(j)$ is the number of those paths that go through vertex i.

5. TOWARDS A SOCIAL NETWORK

In this section we describe how a judgment aggregation problem can be translated into two social networks. This technique was originally introduced in social theory by Breiger [8], although there it was used to analyse the membership of people to groups. We take a slightly different approach and use it in judgment aggregation to obtain graphs that represent agreement between voters and correlation between issues.

5.1 Matrix Translation

In order to obtain social networks from a voting profile we use different matrix transformations. We use the following two types of matrices in this translation.

Definition 3 (Similarity matrix). *Given a profile matrix $\boldsymbol{B}$. The* similarity matrix $\mathcal{B}$ *is obtained from $\boldsymbol{B}$ as follows:*

$$\mathcal{B}_{ij} = \begin{cases} 1 & \text{if } \boldsymbol{B}_{ij} = 1 \\ -1 & \text{if } \boldsymbol{B}_{ij} = 0 \end{cases}$$

Definition 4 (Normalised matrix). *Given a similarity matrix $\mathcal{B}$ of size $n \times m$ and $A = B(B^T)$, where multiplication is ordinary (inner product) matrix multiplication. The* normalised matrix *O of A is constructed as follows:*

$$O_{ij} = \frac{A_{ij} + m}{2}$$

Here we state the main result of this section, showing that it is possible to obtain a matrix from a binary matrix that counts the number of similarities between the rows.

Theorem 1. *Let $\boldsymbol{B}$ be a profile matrix of size $n \times m$, $\mathcal{B}$ the similarity matrix of $\boldsymbol{B}$, and O the corresponding normalised matrix of $\mathcal{B}$. O_{ij} contains the amount of equal elements in row i and j of $\boldsymbol{B}$, i.e.:*

$$O_{ij} = |\{\boldsymbol{B}_{ik} \mid \boldsymbol{B}_{ik} = \boldsymbol{B}_{jk}, 1 \leq k \leq m\}|$$

Proof. Suppose arbitrary rows $\mathbf{B}_i, \mathbf{B}_j$ of some profile matrix $\mathbf{B}$. Let $y = |\{\mathbf{B}_{ik} \mid \mathbf{B}_{ik} = \mathbf{B}_{jk}, 1 \leq k \leq m\}|$ and $x = m - y$. Thus, y is the amount of equal elements between rows i and j in $\mathbf{B}$, and x is the amount of elements that are unequal. Let $\mathcal{B}$ the similarity matrix of $\mathbf{B}$ and O the normalized matrix of $A = \mathcal{B}(\mathcal{B}^T)$ (Definition 4). From the definition of inner product multiplication, it follows that each cell of the matrix A is calculated as follows: $A_{ij} = \sum_{k=1}^{m} \mathcal{B}_{ik}\mathcal{B}_{jk}$. From Definition 3 it follows that if $\mathcal{B}_{ik} = \mathcal{B}_{jk}$, then $\mathcal{B}_{ik}\mathcal{B}_{jk} = 1$, and otherwise $\mathcal{B}_{ik}\mathcal{B}_{jk} = -1$. Thus it follows that $\sum_{k=1}^{m} \mathcal{B}_{ik}\mathcal{B}_{jk} = y - x$. Therefore $A_{ij} = y - x$. Following from applying the normalisation of Definition 4: $O_{ij} = \frac{A_{ij}+m}{2} = \frac{y-x+m}{2} = \frac{y-x+x+y}{2} = y$. □

After transforming a judgment aggregation profile $\mathbf{B}$ into a similarity matrix $\mathcal{B}$, we can perform the following inner product multiplications to obtain two matrices. The first, V^*, measuring the agreement between voters, and the second, I^*, measuring the correlation between the issues.

$$V^* = \mathcal{B}(\mathcal{B}^T) \quad (5)$$

$$I^* = \mathcal{B}^T(\mathcal{B}) \quad (6)$$

We normalise V^* and I^* according to Definition 4, obtaining the the *voter-to-voter matrix V* and the *issue-to-issue matrix I*. It follows now from Theorem 1 that V_{ij} contains the number of equal votes between voter i and j for the same issue. Thus, V_{ij} denotes the number of times that both voters i and j voted "yes" or they both voted "no" for the same issue. Similarly, I_{ij} contains the number of voters that have voted the same for both issues i and j. In this case, this means that I_{ij} denotes the number of voters that have voted "yes" for both issues i and j, or voted "no" for both issues.

Example 5.1 (Continued). *We can translate the matrix $\boldsymbol{B}$ of Example 3.1 that corresponds to this voting profile to a similarity matrix (Figure 1). Next, we calculate V^* and I^* using Eq. (5) and (6). We obtain the the voter-to-voter matrix V and the issue-to-issue matrix I after normalising (Figure 2).*

	p	q	r	z
a	-1	1	1	-1
b	1	-1	-1	-1
c	1	1	1	1
d	1	-1	-1	-1
e	1	-1	1	-1
f	-1	-1	1	-1

Figure 1: Translating a voting profile $\mathbf{B}$

	a	b	c	d	e	f
a	4	1	2	1	2	3
b	1	4	1	4	3	2
c	2	1	4	1	2	1
d	1	4	1	4	3	2
e	2	3	2	3	4	3
f	3	2	1	2	3	4

(a) Voter-to-voter (V)

	p	q	r	z
p	6	2	2	3
q	2	6	4	5
r	2	4	6	3
z	3	5	3	6

(b) Issue-to-issue (I)

Figure 2: Normalised voting matrices

These matrices are both *symmetric* with respect to their main diagonal: If some voter i agrees with a voter j on some issues, then j agrees with i on the same issues as well. This implies *reflexivity*: a voter always agrees with itself over every issue, and similarly for any issue. Therefore, the main diagonal of the voter matrix is always equal to the number of issues, and consequently the main diagonal of the issues matrix is always equal to the number of voters.

5.2 Relational Graphs

The voter-to-voter matrix V and the issue-to-issue matrix I can be represented as two undirected, weighted graphs. In such a graph, a voter (respectively an issue) is represented by a node, and an edge represented the agreement between two voters (resp. correlation between issues). Formally, an edge (i, j) connects two vertices i and j if the matrix entry V_{ij} (resp. I_{ij}) has a value larger than 0. We denote the obtained graphs with $G_V = (V_V, E_V, W_V)$ and $G_I = (V_I, E_I, W_I)$, respectively. Note that the V and I matrices are the same as the weight matrices for the corresponding graphs, i.e. $V = W_V$ and $I = W_I$. We call the graph G_V the *voter graph* and the graph I_V the *issue graph.*

Example 5.2 (Continued). *Figure 3 shows the graphs resulting from the matrices depicted in Figure 2.*

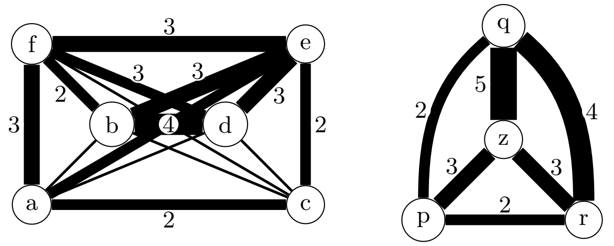

(a) Voter graph (G_V)

(b) Issue graph (G_T)

Figure 3: Voting graphs

Figure 3a shows that the strongest connection is between agents b and d, representing the fact that their ballots are equivalent. Differently, agent c can be considered an outlier due to its weak connections with the other agents. For the issue graph, the issue z seems to be mostly correlated with other issues and therefore has a central position.

For the sake of readability, the edges with a weight of 1 have not been labeled in Figure 3a and reflexive edges have been omitted in both figures.

Due to space limitations, we use the the remainder of the paper to discuss the voter-to-voter matrix and the corresponding voter graph. We discuss further work related to the issue-to-issue matrix and the issue graph in the conclusions (Section 8).

6. THEORETICAL ANALYSIS

In this section we relate the Hamming distance between ballots and profiles to the voter graph, and we discuss a correspondence between the prototype voting rule and the degree centrality measure, introduced in Section 4. We start out with a straightforward equivalence between the Hamming distance between two voters and the edge that connects the two voters in the corresponding voting graph.

Lemma 1. *The Hamming distance between two ballots B_i and B_j is equal to $m - w_{ij}$ in the corresponding voter graph G_V, i.e. $H(B_i, B_j) = m - w_{ij}$.*

Proof. Suppose two ballots B_i and B_j containing m issues. Suppose y to be the amount of issues on which the agents i and j agree. From Theorem 1 it follows that the voter-to-voter normalised matrix V, constructed from a profile $\mathbf{B}$ containing B_i and B_j, has $O_{ij} = y$. Let G_V be the voter graph constructed from V; The weight of the edge between the vertices i and j in G_V is y. The Hamming distance $H(B_i, B_j) = m - y$, hence $H(B_i, B_j) = m - w_{ij}$. □

We use this lemma to obtain an equivalence between the Hamming distance to a profile and the total weight of the corresponding node in the voter graph.

Lemma 2. *The Hamming distance between a ballot B_i and a profile $\boldsymbol{B}$ is equal to $mn - s_i$, where s_i is the sum of the weights of the incident edges of vertex i in the voter graph constructed from $\boldsymbol{B}$:*

$$H(B_i, \boldsymbol{B}) = mn - s_i$$

Proof. Suppose some profile $\mathbf{B}$, a ballot $B_i \in \mathbf{B}$ and a voter graph G_V constructed from $\mathbf{B}$. The Hamming distance between B_i and $\mathbf{B}$ is

$$\begin{aligned} & \sum_{j \in \mathcal{N}} H(B_i, B_j) \\ = & \sum_{j \in \mathcal{N}} m - w_{ij} && \text{(Lemma 1)} \\ = & \; mn - \sum_{j \in \mathcal{N}} w_{ij} \\ = & \; mn - s_i \end{aligned}$$

□

Example 6.1 (Continued). *In Example 3.1, we have $H(a, b) = 3$ and $H(a, \boldsymbol{B}) = 11$. In the corresponding graph in Figure 3a we have that $w_{ab} = 1$ and thus $m - w_{ab} = 4 - 1 = 3$, which corresponds to the Hamming distance between a and b. Moreover, $s_a = 13$ (including the reflexive weight of 4), so $mn - s_a = 24 - 13 = 11$, which corresponds to the Hamming distance between a and the profile $\boldsymbol{B}$.*

Since the prototype voter rule selects the voter that minimizes the distance with the profile, we can obtain the following equivalence:

Lemma 3. *The prototype rule PRO (Definition 2) selects the voters corresponding to the maximum total weight vertices in the voter graph, i.e.:*

$$PRO(\boldsymbol{B}) = \underset{i \in V_V}{\operatorname{argmax}}\, s_i.$$

Proof.

$$\begin{aligned} PRO(\mathbf{B}) &= \underset{B \in \text{Supp}(B)}{\operatorname{argmin}} \; \mathcal{H}(B, \mathbf{B}) && \text{(Definition 2)} \\ &= \underset{i \in V_V}{\operatorname{argmin}}(mn - s_i) && \text{(Lemma 2)} \\ &= \underset{i \in V_V}{\operatorname{argmax}}\, s_i \end{aligned}$$

□

Next, we obtain that the average voter rule corresponds to the node with the highest degree centrality when the tuning parameter $\alpha = 1$:

Theorem 2. *The prototype rule select those individual ballots that have the maximal degree centrality value when $\alpha = 1$. Suppose $\alpha = 1$:*

$$PRO(\boldsymbol{B}) = \underset{i \in V_V}{\operatorname{argmax}}\, C_D^{W\alpha}(i)$$

Proof. Follows directly from Eq.(1) and Lemma 3 □

	1	2	3	4	5
a	0	1	1	1	1
b	1	0	0	0	0
c	0	1	1	0	0
d	0	0	0	1	1
Maj:	0	?	?	?	?
AVR:	0	1	1	1	1
	0	1	1	0	0
	0	0	0	1	1

(a) Profile

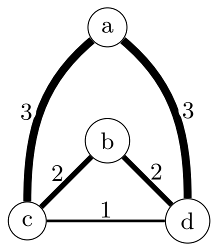

(b) Voter graph

Figure 4: Judgment aggregation example

6.1 Varying the tuning parameter

We have shown a correspondence between the the prototype voter and the degree centrality when the tuning parameter $\alpha = 1$, but we are interested in the effect for varying α as well. Consider the voting scenario together with the majority outcomes and the prototype voter outcomes that are depicted in Figure 4a and the corresponding voter graph in Figure 4b. The outcome of the degree centrality for varying α are depicted in Figure 5. As can be seen from the table, for $\alpha = 1$, the degree centrality measure corresponds to the s_i measure, so node a, c and d are all chosen as the most representative voter. When $\alpha < 1$ the amount of connections play are larger role and only c and d are chosen as the winner, while for $\alpha > 1$ the weight of the edges play are larger role and a is picked as the winner. Thus, the intuition seems to be that in some cases, using the degree centrality we are able to obtained a more fine-grained approach than using the prototype voter rule.

It is interesting to compare these outcomes with the vector of "average votes" $(\frac{1}{4}, \frac{2}{4}, \frac{2}{4}, \frac{2}{4}, \frac{2}{4})$, showing for each issue the proportion of voters who chose 1 rather than 0. This demonstrates that only the first issue is relatively uncritical, but that it is not trivial to decide on the other four issues. This is in line with the outcomes of both the majority voter and the prototype voter, but as we have shown the degree centrality is able to give a more specific outcome. It thus seems the structure of the graph can be exploited to fine-tune distance-based outcomes. In the next section we will try to strengthen these intuitions by performing several experiments.

Vertex	s_i	$C_D^{W\alpha}$ when $\alpha =$			
		0	0.5	1	1.5
a	**4**	2	2.5	**4**	**6.3**
b	2	2	2	2	2
c	**4**	**3**	**2.9**	**4**	5.5
d	**4**	**3**	**2.9**	**4**	5.5

Figure 5: Degree centrality scores when different values of α are used.

7. EMPIRICAL ANALYSIS

To evaluate the outcomes of the centrality measures, we provide an empirical analysis in this section. We compare the degree, closeness, and betweenness measures with the majority and the prototype rule.

7.1 Experimental Setup

The setup of the experiment that we have performed[1] consists of a judgment aggregation problem with n voters and m issues, with the integrity constraint $IC = p_1 \wedge p_2 \wedge \ldots p_{m-1} \leftrightarrow p_m$. The votes are generated pseudo-randomly such that all votes are complete and the generated ballots are consistent with the constraint.

In order to compare the different measures we use the distance-based rule (Definition 1) as the *base measure*. We compare the outcome of each measure with the base measure using the Hamming Distance.

If the distance-based rule produces multiple outcomes, we consider the distance between a measure and the base rule to be the minimum distance between the result of the measure and the set of results of the base rule. For instance, if the base measure outcome is $\{[0,1,1],[1,1,1]\}$ and the majority outcome is $\{[0,1,0]\}$, then the distance between the two outcomes is 1. If a measure produces multiple outcomes, we measure the distance to the base measure for each result. All these distances are stored in a list L_M for each measure M.

The experiment is repeated for E times, after which each list L_M is used to compute the mean, the standard deviation σ and the average number of outcomes per benchmark O_{avg}, i.e. $O_{avg} = \frac{|L_M|}{E}$ for the measure M. The value O_{avg} can be seen as a measure for resoluteness: The closer this number is to 1, the more resolute the voting rule is. Additionally, for the majority measure we also compute the ratio of inconsistent outcomes, which we denote with $\perp$.

7.2 Results

Since the centrality measures used in this paper are based on graph theory, these measures tend to produce interesting results on large graphs due to the more dependencies and similarities between the agents. Initial tests have shown that the centrality measures produce the best results when the tuning parameter $\alpha = [1, 4]$.

Figure 6 shows an experiment with a relatively large group of 25-31 voters and 5-10 issues. We use ranges for the number of voters and issues to avoid that the outcome depends on the structure of the voting problem (for instance, if the number of voters is uneven, the majority rule will always be resolute). From this figure it can be seen that, as shown in Section 6, the prototype rule corresponds the degree centrality measure with the tuning parameter $\alpha = 1$. When the value of α increases from 1 to 4, the measure becomes slightly more resolute but also less precise. While resoluteness increases with about 4.5%, the mean distance increases with about 47%. So the distance increases about ten times as much as the resoluteness. This leads us to the conclusion that increasing α does not lead to better results for the centrality degree.

When we consider the closeness centrality measure, we see that this measure becomes more precise when the value of α increases, while retaining resoluteness. In fact, when α increases from 1 to 4, the mean decreases with about 48%, while the resoluteness stays constant. In the end, the measure is only slightly less precise than the prototype voter and much more resolute. Intuitively, it makes sense that

[1]The experiment has been coded in Java and can be found on the web, but we have left it out in this version to preserve anonymity of the authors. They will appear in the final version.

closeness performs well for higher α, because few edges with lower weights are more likely to correspond to similarity in voting than more edges with higher weights, simply because increasing the edges between two nodes makes it less likely that the weights still correspond to similarities of the two nodes. For instance, if two nodes are connected by an edge with weight 1, it is sure that the corresponding agents have voted the same for one issue, but if there are two edges with a weight of 2 between these nodes, there is no guarantee that the agents corresponding to these nodes have reached any agreement because this depends on the middle agent as well.

Lastly, we point out that the betweenness measure does not seem to perform very well. This can be explained by the fact that a great proportion of nodes in a network generally does not lie on a shortest path between any two other nodes, and therefore receives the same score of 0 [21].

To study the resoluteness of the closeness measure in isolation, we have compared it to the majority measure while varying the number of agents and the number of issues in the range [2-10]. The results of this comparison are depicted in Figure 7. The left chart contains the average number of outcomes in the case of the majority measure. As can be seen, the majority rule tends to be irresolute when the number of voters is even, and the number of outcomes seem to increase exponentially with the number of issues. The right chart shows that the amount of outcomes of the closeness centrality measure is consistently lower.

From our empirical results we can confirm that the degree centrality measure corresponds to the PRO voter when $\alpha = 1$. Moreover, in the case of closeness centrality, increasing α results in less outcomes, while the outcomes remain generally acceptable. Lastly, the betweenness centrality seems unsuitable as a voting rule.

Voting rule	α	mean	σ	O_{avg}	$\perp$
Majority	-	0.03	0.17	1.43	0.03
PRO	-	0.17	0.39	1.38	-
Degree	1	0.17	0.39	1.38	-
	2	0.21	0.45	1.34	-
	3	0.24	0.51	1.32	-
	4	0.25	0.53	1.32	-
Closeness	1	0.44	0.62	1.02	-
	2	0.28	0.5	1.04	-
	3	0.24	0.46	1.05	-
	4	0.23	0.48	1.02	-
Betweenness	1	2.09	1.15	10.13	-
	2	2.05	1.2	4.16	-
	3	2.25	1.26	2.13	-
	4	2.58	1.27	1.86	-

Figure 6: Parameters: $E = 5000$, $n \in \{25, \ldots, 31\}$, $m \in \{5, \ldots, 10\}$

7.3 Complexity Analysis

The time complexity of computing the voter matrix using the standard schoolbook multiplication is $O(n^2m)$. The time complexity of computing the degree is $O(n^2)$, and that of the closeness and betweenness is $O(nm+n^2 \log n)$ in both cases [20] [7].

Therefore the time complexity to compute the outcome of a judgment aggregation problem using the centrality measures is dominated by the matrix multiplication in all cases, resulting in a time complexity of $O(n^2m)$, which roughly corresponds to a cubic time complexity of $O(n^3)$, assuming that the number of issues is not much larger than the number of voters, or at least in the same magnitude.

8. CONCLUSIONS AND FUTURE WORK

To the best of our knowledge, the present paper is the first attempt to model the (social) relations between judgment aggregation using techniques from social network analysis. We have reformulated a classical technique that analyses the membership of people to groups [8] and applied them to judgment aggregation, which allows us to model agreement between agents and correlation between voters.

We have studied the agreement between agents as a social graph, and showed that well-known notions in judgment aggregation such as Hamming distance and the prototype voter can be equivalently formulated using the graph.

We empirically analysed how degree, closeness and betweenness centrality measures perform with respect to the distance-based rule and the prototype voter rule. From the results it can be noticed that the degree centrality measure approximates quite well the prototype voting rule, the closeness centrality measure turns out to be more resolute but a little less accurate when using a large enough tuning parameter α, and the betweenness centrality measure is unsuited to be used to solve judgment aggregation problems.

Although we have provided both theoretical and empirical results, much remains to be done. We have restricted the theoretical analysis to the degree centrality with the tuning parameter $\alpha = 1$, but we plan to investigate whether similar results can be obtained for the other two measures, closeness and betweenness, as well.

We have focused on the relations that exist between the agents in a judgment aggregation scenario, but as we have seen in Section 5.2, it is possible to perform a similar analysis on the issues as well. We deem interesting to study whether the emerging relations between the issues are correlated to the given integrity constraints, and whether some relations between the issues may emerge even when no integrity constraints were given (i.e.: when the correlation between two issues is not common knowledge but some of the agents are aware of it).

Last but not least, depending on the findings concerning the importance of relations between the agents and the issues, we also aim at providing additional judgment aggregation rules exploiting the emerging relations in a profile.

9. REFERENCES

[1] K. J. Arrow. *Social Choice and Individual Values.* John Wiley and Sons, New York, NY, 1951.

[2] I. Asimov. *Robot Dreams.* Berkley, 1986.

[3] A. Barrat, M. Barthélemy, R. Pastor-Satorras, and A. Vespignani. The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences of the United States of America*, 101(11):3747–3752, Mar. 2004.

[4] R. Booth and M. Caminada. Quantifying disagreement in argument-based reasoning. *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, 2012.

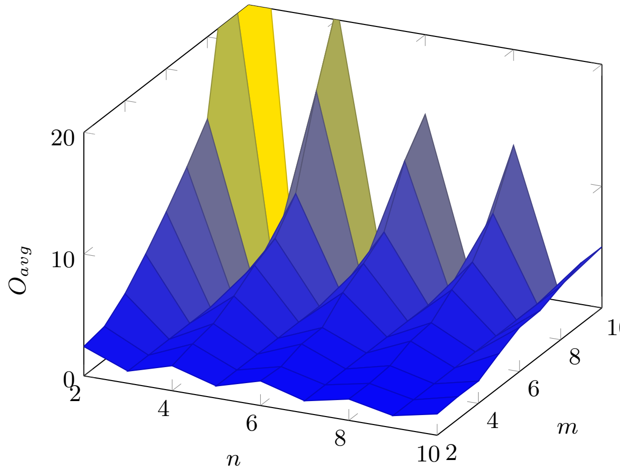

(a) Majority measure

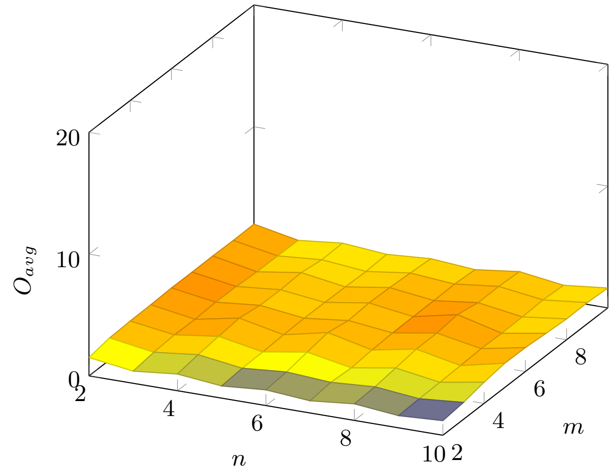

(b) Closeness centrality measure ($\alpha = 4.0$)

Figure 7: Average number of outcomes (z-axis) for varying voters (x-axis) and issues (y-axis)

[5] C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. Cp-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

[6] U. Brandes. A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25:163–177, 2001.

[7] U. Brandes. A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25:163–177, 2001.

[8] R. Breiger. The duality of persons and groups. *Social forces*, 1974.

[9] M. W. A. Caminada and G. Pigozzi. On judgment aggregation in abstract argumentation. *JAAMAS special issue on Computational Social Choice*, 22:64–102, 2011. (Online First published on November 3, 2009).

[10] A. D'Andrea, F. Ferri, and P. Grifoni. An overview of methods for virtual social networks analysis. In *Computational Social Network Analysis*, Computer Communications and Networks, chapter 1, pages 3–25. Springer London, London, 2010.

[11] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, 1959.

[12] U. Endriss and U. Grandi. Graph aggregation. In *Proceedings of the 4th International Workshop on Computational Social Choice (COMSOC-2012)*, September 2012.

[13] L. C. Freeman. Centrality in social networks conceptual clarification. *Social Networks*, page 215, 1978.

[14] U. Grandi and U. Endriss. Binary aggregation with integrity constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, July 2011.

[15] U. Grandi and G. Pigozzi. On compatible multi-issue group decisions. In *Proceedings of the 10th Conference on Logic and the Foundations of Game and Decision Theory*, 2012.

[16] C. List and B. Polak. Introduction to judgment aggregation. *Journal of Economic Theory*, 145(2):441–466, Mar. 2010.

[17] C. List and C. Puppe. Judgment aggregation: A survey. *Handbook of Rational and Social Choice*, 2009.

[18] M. K. Miller and D. N. Osherson. Methods for distance-based judgment aggregation. *Social Choice and Welfare*, 32(4):575–601, 2009.

[19] M. E. J. Newman. Scientific collaboration networks. II. shortest paths, weighted networks, and centrality. *Physical Review E*, 64(1), June 2001.

[20] K. Okamoto, W. Chen, and X.-Y. Li. Ranking of closeness centrality for large-scale social networks. In *Proceedings of the 2nd annual international workshop on Frontiers in Algorithmics*, FAW '08, pages 186–195, Berlin, Heidelberg, 2008. Springer-Verlag.

[21] T. Opsahl, F. Agneessens, and J. Skvoretz. Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks*, 32(3):245–251, July 2010.

[22] C. Pinheiro. *Social Network Analysis in Telecommunications*. Wiley and SAS Business Series. Wiley, 2011.

[23] A. Salehi-Abari and C. Boutilier. Empathetic social choice on social networks. In *Fourth International Workshop on Computational Social Choice*, 2012.

[24] M. Slavkovik. *Judgment aggregation for multiagent systems*. PhD thesis, Univerity of Luxembourg, 2012.