HAL open science

Edge-Compressed Majority Graph: Where Social Choice Meets Information Visualization

Nikos Karanikolas, Renaud Blanch, Sylvain Bouveret

▶ **To cite this version:**

Nikos Karanikolas, Renaud Blanch, Sylvain Bouveret. Edge-Compressed Majority Graph: Where Social Choice Meets Information Visualization. Sixth International Workshop on Computational Social Choice, Jun 2016, Toulouse, France. hal-01399828

HAL Id: hal-01399828

https://hal.univ-grenoble-alpes.fr/hal-01399828

Submitted on 21 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Edge-Compressed Majority Graph: Where Social Choice Meets Information Visualization

Nikos Karanikolas, Renaud Blanch and Sylvain Bouveret

Abstract

Collective decisions are everywhere: choosing central or local governments, selecting a candidate to hire for an open position, choosing a restaurant to share a dinner with some friends are examples of collective decision making situations. Social Choice provides a lot of methods which can help people making a decision in such situations. However, the diversity of these voting procedures and the mathematical background necessary to understand them can be seen as obstacles to the use of these methods in everyday situations by laypersons. We claim that information visualization techniques can help a lot the democratization of social choice, by providing people with some easily interpretable information and, in the end, helping them making informed collective decisions. In this paper, we present the Edge-Compressed Majority Graph, a technique dedicated to the visualization of the majority graph of a preference profile. Using an insight-based evaluation method, we show that this technique gives better results in conveying information about the preferences than other classical visualization techniques.

1 Introduction

The need for collective decisions has played an important role from the beginning of mankind's first organized societies till nowadays. Democracy, which is considered from an ethical point of view the ideal form of government is based on elections, where all citizens have an equal say in the decisions that affect their lives. A major problem that arises in democratic societies is how to interpret the will of people and bring out a satisfying collective decision. Social Choice, and especially Voting Theory, is the field that addresses the problem by providing a set of methods aiming at aggregating a set of individual preferences into a collective preference or decision. The paradigmatic problem is here to aggregate a set of N rankings representing the voters preferences on a set of m candidates into a collective ranking using a *voting rule*.

Throughout the history of social choice many rules have been proposed, and each one of them is trying to reflect the socially fairest outcome. In the case where there are only two candidates (and an odd number of voters), majority voting[1] is unanimously considered a perfect method of selecting the winner. However, when there are three candidates or more, no such obvious rule choice exist, and there exists numerous voting methods that all have different axiomatic properties. Among these methods, a prominent family of rules are based on the Condorcet principle. This principle was introduced by Nicolas de Condorcet, the founding father of the mathematical theory of voting, who suggested a rule that extends majority voting to multiple candidates [9]. A candidate x is said to beat candidate y in a *pairwise election (comparison)* if the majority of voters prefer x to y, *i.e.*, rank x above y. A candidate that beats every other candidate in a pairwise election is the winner of the entire election. Such a *Condorcet winner*, when it exists, is generally considered a good social consensus. The family of *Condorcet-consistent* voting rules aim at reconciling the

[1] In majority voting, each voter only votes for one candidate (her top one). The candidate receiving the highest number of votes wins the election.

Condorcet winner with the fact that it may not exist, by electing it when it exists. This family includes, among others, *Simpson's rule*, *Copeland's method*, [8], *Dodgson's rule* [4] and *Young's method* [21]. Some of these methods are based on the *(weighted) majority graph*: a directed graph where each node represents a candidate and the arc (x, y) exists if and only if x beats y in a pairwise comparison. Even if some Condorcet-consistent methods are not directly based on it, the majority graph is a good way to understand the rationale behind these methods (namely, pairwise comparisons).

A fundamental question of Social Choice is the choice of the voting rule to use when there are three candidates or more, as different voting rules can yield very different collective preferences (and hence different winners). The traditional approach in social choice theory is prescriptive: facing a collective decision making problem, a society will choose a voting procedure that satisfies good axiomatic properties (monotonicity, reinforcement...). However, making an informed choice about the voting method requires a strong mathematical background that most people do not have, leaving them with no solution but to choose a black-box voting method by default. We believe that a good alternative is to use (graphical) information visualization techniques to help people understanding the structure of preferences without prescribing any ranking given by a voting method.

The domain of *Information Visualization (InfoVis)* deals with designing visual representations of abstract information to help the user increasing her knowledge about the internal structure of the data and causal relationships in it. Information Visualization aims at providing techniques for the users to make discoveries, decisions, or explanations about patterns, groups of items, or individual items. Most works in InfoVis aim at providing visualization techniques that do not require any strong mathematical background or knowledge from the user for her to be able to interpret the information. Our research direction is to try to apply these techniques on social choice problems, where the information we want to visualize is the votes and the interpretation of these votes.

As a first step, we focus, in this work, on the visualization of the weighted majority graph. We believe that the weighted majority graph is an interesting intermediate aggregation level between the complete preference profile and the collective preference issued by a particular voting rule. It already contains compressed information about the individual preferences but can still act as a proxy for a lot of Condorcet-consistent voting rules. We also believe that a decision maker can learn a lot of information about the structure of individual and collective preferences from an efficient graphical representation of the weighted majority graph. It should be noticed, though, that choosing a representation based on this graph is not neutral. It orients the decision towards candidates that are good according to the Condorcet principle. This is not without loss of generality and will favour some types of candidates.

Giving an efficient graphical representation of the majority graph is not straightforward. Being a complete directed graph, the traditional node-link representation quickly becomes unusable as the number of candidates grows. In this paper, we propose (in Section 4) a visualization called the Edge-Compressed Majority Graph, and compare it with the classical node-link representation and a representation based on the adjacency matrix. The experiments (provided in Section 5) show that people are able to extract a lot more useful pieces of information from the Edge-Compressed Majority Graph visualization technique than from the other classical ones.

2 Related work

The application of Information Visualization in Social Choice is a relatively new field, however visualization in elections is widely used in many forms. Occasionally, there are many graphs and animations published in the newspapers or news websites visualizing the results of different kinds of political elections. However, most voting rules used in political contexts do not ask the voters to give their preferences in the form of linear orders (rankings), but rather ask e.g. either to give one's preferred candidate, or to approve a set of candidates. This is why techniques based on the majority graph (like Condorcet consistent rules and the visualization techniques we use in this paper) are not relevant in this context. This motivated us to study meticulously the visualization of more complex election settings and voting rules, like the one promoted by Social Choice theorists.

To the best of our knowledge, the study of visualization techniques applied to Condorcet-consistent voting rules is new. However, the usage of majority graphs in this kind of voting rules have been being studied for years in the Social Choice literature, with the application of tournament solutions. One such method, which was proposed by Copeland [8], constitutes one of the most important tournament solutions. Other significant tournament solutions were proposed by Levchenkov [15, 13]; Fishburn [12] and Miller [16]; and Banks [1]. An analytic survey of tournament solutions can be found in Laslier's book [14].

Perhaps the closest work to our approach is the one by Betzler et al. [3]. In their work, they propose a set of data reduction techniques to compute the winner of the Condorcet-consistent Kemeny's rule (which is NP-hard to compute). One of their data reduction techniques is based on the computation of clusters of candidates using the strongly connected components in the majority graph. As we will see in Section 4, our Edge Compression Majority Graph uses a very similar reduction technique. The main difference is that we use it for visualization purposes, where Betzler et al. use it for theoretical (complexity) results.

3 Preliminaries

We consider a set $N = \{1, \ldots, n\}$ of *voters* and a set C of m *candidates*. Each voter i has a linear order $\succ_i$ on C, where $x \succ_i y$ means that voter i prefers candidate x to y. Let $\mathcal{L}_C$ be the set of linear preferences over C. A *preference profile* $\succ = \langle \succ_1, \ldots, \succ_n \rangle \in \mathcal{L}_C^n$ is a collection of preferences for all the voters. A *voting correspondence* is a mapping $f : \mathcal{L}_C^n \to 2^C \setminus \{\emptyset\}$ from preference profiles to non empty subsets of candidates, the co-winners of the election. A *voting rule* maps a preference profile to a unique winner. Most voting rules are built from the composition of a voting correspondence and a tie-breaking rule.

Let $\succ \in \mathcal{L}_C^n$. The *majority margin* m is the function that maps each pair (x, y) of candidates to the difference between the number of voters that prefer x to y and the number of voters that prefer y to x, namely: $m(x, y) = |\{i \in N : \ x \succ_i y\}| - |\{i \in N : \ y \succ_i x\}|$. The *majority relation* is the binary relation $\succeq_{\text{Maj}}$ defined as the subset of elements $(x, y) \in N^2$ such that $m(x, y) \geq 0$. The binary relation $\succ_{\text{Maj}}$ is the strict part of $\succeq_{\text{Maj}}$; therefore if $x \succ_{\text{Maj}} y$, we say that x (strictly) beats y in a pairwise election (or comparison). The *majority graph* is the directed graph representing $\succeq_{\text{Maj}}$. A *Condorcet winner* is a candidate that strictly beats every other candidate in a pairwise election.

Among the set of Condorcet-consistent voting rules, some of them are based on the majority relation. In particular, the Copeland rule works as follows. For each candidate x, the Copeland score of x is the number of candidates y such that $x \succ_{\text{Maj}} y$. The candidate with the highest Copeland score is the winner of the election. Some other Condorcet-consistent voting rules, among which the Kemeny rule, are based on the information provided in the weighted majority graph. Finally, some rules use the information contained in the full preference profile. In particular, the Dodgson rule works as follows. For each candidate x, the Dodgson score of x is the smallest number of sequential exchanges of adjacent candidates in the preference profile needed to make x a Condorcet winner.

In the following, we will have to resort to a Condorcet consistent voting rule when an order is required for some specific graphical attribute (as is the case in the Pairwise Comparison Matrix and our Edge-Compressed Majority Graph techniques. Dodgson's voting rule is intuitively appealing and one of the most well-studied voting rules in the literature. However it is hard to compute since the Dodgson score decision problem is NP-complete, as shown by Bartholdi, Tovey and Trick [2]. In order to apply a Condorcet-consistent voting rule that will be computed efficiently (polynomial time) we will use the one proposed by Caragiannis et al. [7]. They proposed the following Dodgson approximation rule which is an extension of Tideman's simplified Dodgson rule [20, pages 199-201]. Consider a profile $\succ \in \mathcal{L}_C^n$ and a candidate x. If a candidate x is a Condorcet winner, then it has Extended Tideman's Simplified Dodgson (ETSD) score $\mathrm{sc}_{\mathrm{Td}'}(x, \succ) = 0$. Otherwise:

$$\mathrm{sc}_{\mathrm{Td}'}(x, \succ) \quad = \quad m \cdot \mathrm{sc}_{\mathrm{Td}}(x, \succ) + m(\log m + 1), \tag{1}$$

where $\mathrm{sc}_{\mathrm{Td}}(x, \succ)$ is simplified Dodgson score of candidate x defined as:

$$\mathrm{sc}_{\mathrm{Td}}(x, \succ) \quad = \quad \sum_{y \in A \setminus \{x\}} \max\{0, n - 2 \cdot |\{i \in N : x \succ_i y\}|\}. \tag{2}$$

The candidate with the minimum score wins.

4 Visualizing the majority graph

In this section we describe how we graphically visualize weighted majority graphs. We consider three types of visualization techniques: the Pairwise Comparison Matrix, the Weighted Majority Graph and the Edge-Compressed Majority Graph. All of these techniques aim at extracting a synthetic information from a voting situation in order to allow a decision maker understand the structure of the voting profile. Clearly, the general purpose of our work is to provide a set of visual representations of the voting data which are clearly understandable without prior knowledge on any voting method. We also emphasize that all of our methods can be employed for the visualization of any kind of weighted tournament graph, whether they represent the weighted majority graph of a voting situation or not.

Before introducing our visualization techniques, we should note that ordering the candidates is a crucial aspect. This is inherent to graphical visualization, since most graphical attributes naturally induce an ordering (*e.g.* left-to-right for abscissa, top-to-bottom for ordinate, green-to-red for colors). To deal with this aspect, we had two obvious choices. Either randomly map values of graphical attributes (x, y, color...) to candidates or choose a specific ordering and use graphical attributes to convey it. We chose the second approach for the sake of graphical clarity (randomness yield representations that are harder to read and interpret). We have used a Condorcet-consistent aggregation method, more precisely ETSD scores ($\mathrm{sc}_{\mathrm{Td}'}$) (because of its similarity to the Dodgson's), as input order to the graphical representations.[2] We know that this is at the price of inducing a strong bias in the way users perceive the visualization.

4.1 Pairwise Comparison Matrix

The first visualization technique is the *Pairwise Comparison Matrix (PCM)*, which simply displays the majority margin between each pair of candidates using a matrix of colors. Every entry of the matrix $a_{i,j}$ represents the majority margin $m(i, j)$, with $i, j \in [\![1, m]\!]$, encoded

[2]Note that using Tideman's would lead to the same ranking as the Td'. We prefer to use the latter rule because it is a Dodgson-approximation and thus preserves the characteristics of Dodgson's rule (*e.g.*, the proximity of candidates to becoming Condorcet winners). Td' has the advantage over Dodgson's that it is polynomial and satisfies a lot of well-known social choice properties, like monotonicity and homogeneity.

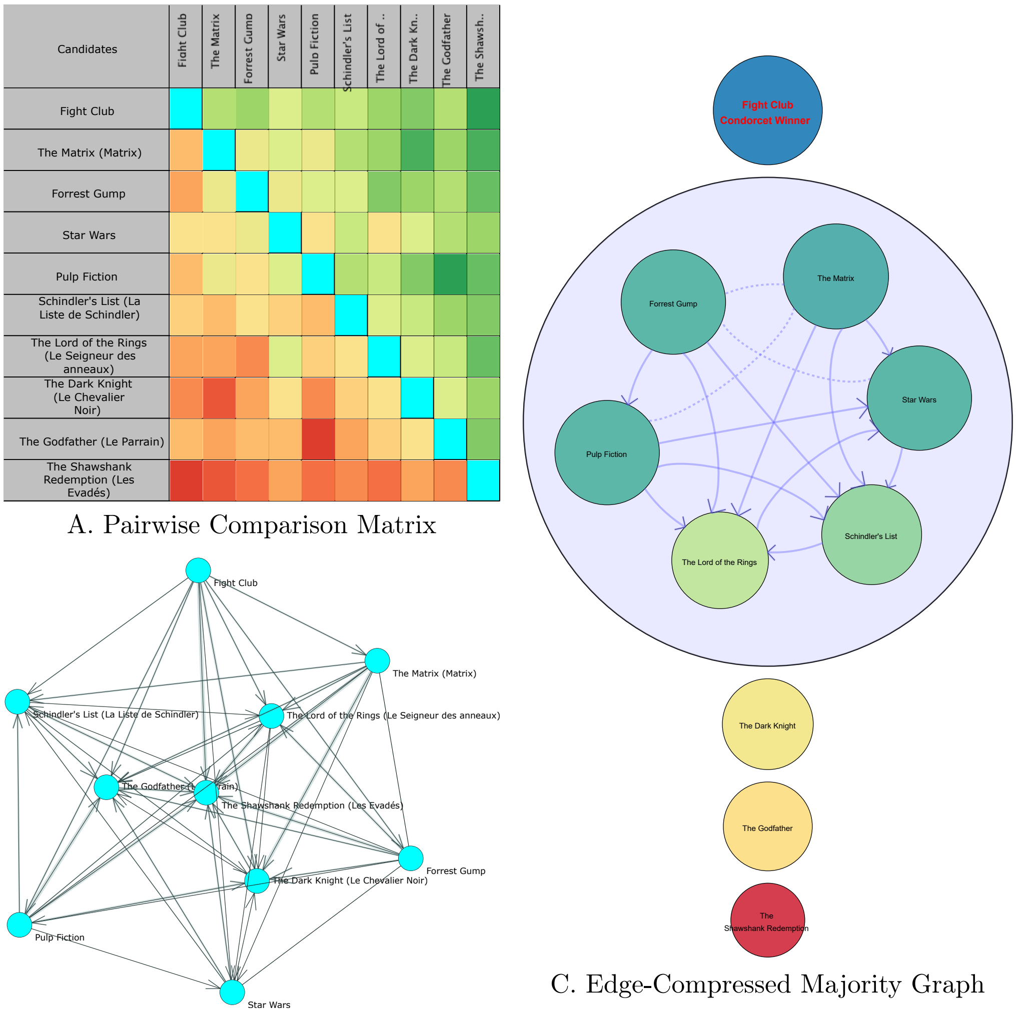

A. Pairwise Comparison Matrix

B. Weighted Majority Graph

C. Edge-Compressed Majority Graph

Figure 1: Different visualizations for the movie poll

using a green to red color scale. More precisely, we use the "diverging category with ten data classes" (10-class RdYlGn) scale [6] and map the values (percentage) to these colors. When $m(i,j) > 0$ (resp. < 0), $a_{i,j}$ is a variant of green (resp. red) color where the saturation level is simply $m(i,j)/n$ (resp. $m(j,i)/n$). If $m(i,j) = 0$ the $a_{i,j}$ cell is a yellowish color. The rows and columns are ordered according to the ETSD score of the corresponding candidate. An example of PCM is shown in Figure 1.A. The data set used in this figure is the movie poll data set described in Section 5.1.

4.2 Weighted Majority Graph

The next visualization technique is the *Weighted Majority Graph (WMG)*, which is a simple node-link representation of the majority graph with directed arcs for the strict part of the majority relation and undirected edges when two candidates tie. The weight of each arc (x, y) equals to $m(x, y)$ and thus the sum of weights of the incoming arcs for a node x equals the simplified Dodgon score of candidate x. We use the width of each arc (x, y) as a graphical attribute encoding the weight of (x, y). Moreover, the graphical layout displays the candidate(s) who have the minimum simplified Dodgson score at the top. An example of the WMG for the aforementioned movie poll can be seen in Figure 1.B.

4.3 Edge-Compressed Majority graph

The last visualization technique, which is the one we propose, is the *Edge-Compressed Majority Graph (ECMG).* Our objective is to improve upon the WMG by providing a visualization where the most valuable information —namely, comparisons between candidates and *voting cycles*— can be easily extracted. Our method distinguishes the candidates that can be linearly ordered from the ones that are involved in cycles. The novelty of our technique is that it only shows the edges between the candidates of the latter type. In order to do that, we combine edge compression techniques in the visualization of dense graphs [10] with the use of the axis y as a graphical attribute encoding a linear ordering. Moreover, we rely on concepts specific to Voting Theory, and especially on the properties of the Smith set [19], to perform the edge compression. The *Smith set* is the smallest non-empty set of candidates in a particular election such that each member of this set beats every other candidate outside the set in a pairwise comparison.

The edge compression algorithm works as follows. Starting from the whole set of candidates C, we compute a Smith set, and append it to the list of clusters. Then we reapply the same procedure iteratively on the remaining majority graph until no candidate remains. We end up with a sequence of clusters formed by the successively computed Smith sets. These clusters have the nice property that every node inside the same cluster have exactly the same set of arcs with nodes outside the cluster. This set of arcs can thus be "compressed" using arcs between clusters.

The algorithm we use for the computation of a Smith set is an adaptation from the algorithm presented in the Wiki for Election Methods [11], that is an adjusted version of the *Floyd-Warshall* algorithm. We can observe that we could have used a different graph-theoretic perspective, because our technique can be seen equivalent to a decomposition into the strongly connected components of a directed graph. This decomposition can produce the same clusters of candidates and then we can apply a topological ordering on the components to get the same results as in our technique. Note that the Smith sets found during our technique are exactly the same as the strongly connected components of the graph given the following condition about ties: when there is a tie between two candidates then the undirected edge must be replaced by 2 directed edges heading towards both directions. The difference in time complexity for both techniques is not remarkable, so we preferred the Smith set decomposition because the implementation is simpler and relies also on notions of voting theory. Note that the computation of the Smith set needs time $\mathcal{O}(m^3)$ using the adaptation of the Floyd-Warshall algorithm and thus globally runs in $\mathcal{O}(m^4)$. We could have used the Kosaraju algorithm for computing the Smith set and thus reducing the time complexity of our technique to $\mathcal{O}(m^3)$ but we preferred a simpler implementation and the complexity does not really matter as the number of candidates (m) is in general low.

Graphical representation of the cluster graph At this point, we have to represent graphically a sequence of clusters —each cluster being a subgraph of the majority graph. The intuitive meaning of this sequence is that each candidate in a cluster is better than all the other candidates in the subsequent clusters. A natural way of representing linear orders is to use the abscissa or the ordinate as the graphical attribute encoding the rank of each element in the order. In our visualization, we use the latter, and the clusters are displayed on top of each other from the first to the last one. Doing this, we do not need to draw any arc between the clusters as the relation between two clusters can be naturally deduced from their respective ordinates.

Each cluster is displayed as a circle (or rectangle) containing the subgraph of the majority graph. If a cluster contains only one candidate we do not draw the border of this cluster but directly draw the candidate itself.

Each candidate (inside a cluster) is drawn as a circle, whose size and color depend on the candidate's ETSD score. Note that we use the minimum and maximum scores of the whole set of candidates and not only the ones belonging to the cluster. We do that in order to have a global visual representation of the candidates, otherwise there might be a case where a losing candidate that belongs to a lower ranked cluster has bigger size than a candidate who beats her. We use the same voting rule as above because, as already mentioned, it is easily computed and also produces less ties among the candidates compared to other Condorcet-consistent rules, like the Copeland's rule. Imagine the case where there are multiple members of a Smith set, then the Copeland's method often yields ties. For example, if there is a three-candidates voting cycle, each candidate will have exactly one loss, and there will be an unresolved tie between the three and there is no winner declared. On the other hand ETSD procedure computes the score for each candidate by adding the majority margins which is more unlikely to be the same for two candidates. Therefore, this rule is important in order to rank the nodes inside the cluster (of nodes) composing the Smith set and thus have a better visual appearance. In order to compute the circle's diameter for a node we use the following procedure: we assign maximum and minimum values of diameter that correspond to the candidates having the minimum and maximum scores accordingly, and then the node's diameter is calculated proportionately according to these values. The node's color is computed from the score. First, we choose a palette of 8 colors that extends from blue color to red (the palette used is the diverging 8-class spectral [6]). Blue color is assigned to the candidate with the minimum score while red color to the candidate with the maximum score. Then, we compute the difference between the maximum and minimum scores of the candidates and divide this distance in 7 intervals of equal size, where the endpoints of the intervals correspond to each one of the 8 colors. Every candidate is assigned then to the appropriate interval according to her score. The color of the candidate is computed by the linear interpolation in the RGB color space of the endpoints of the interval she belongs to.

After computing all nodes' attributes, our algorithm checks the size of the Smith set. When the set is a singleton then the node corresponding to the sole candidate of the Smith set is drawn with the computed diameter and color. Otherwise, the members of the Smith set form a cluster of nodes. When there is a total (linear) order on the members of the Smith set[3] then the nodes are drawn inside a rectangle in a vertical arrangement. The candidate of the set with the minimum score is ranked on top of the rectangle and the other candidates follow her according to their score.[4] If two or more candidates have the same score, then they are drawn at the same ordinate. Now, when there is not a total order on the members of the Smith set then voting cycles between candidates occur and the nodes are drawn inside a circle. The nodes inside the circle are drawn according to their candidate's score, where a candidate is drawn on a higher vertical position compared to another candidate if the former has score less or equal to the latter one.

The final step of our technique is edge drawing. Any edge of the graph represents the result of the pairwise comparison between two candidates. A directed edge is drawn as an arrow from candidate x to candidate y if and only if x beats y. If the result of the pairwise comparison between two candidates is a tie then an undirected edge is drawn as a dotted curve (or line). The edges (directed or undirected) are drawn using cubic Bézier curves,

[3] Recall that the order of the members in the Smith set is done according to their score.

[4] We use ETSD rule to compute the score.

circular arcs or straight lines. The edge type depends on the position of the nodes as we try to find the best visual way to represent them. Specifically, the type of an edge is drawn according to the following guidelines, depending on whether the encompassing cluster is circular (i) or rectangular (ii). Case (i): If the nodes in the graph are close to each other then the edge is a circular arc that follows the perimeter of the cluster's circle, otherwise it is a cubic Bézier curve. Case (ii): If the nodes in the graph are close to each other then the edge is a straight line, otherwise it is a cubic Bézier curve.

The algorithm for the computation of the ECMG visualization is presented in Appendix (Algorithm 1). This algorithm resorts to procedure ComputeSmithSet for computing the Smith set on a given set of candidates, and to Algorithm 2 to draw the clusters of nodes (successive Smith sets). This latter algorithm in turn resorts to two procedures DrawNode and SortNode which are respectively in charge of drawing a circular node of a given diameter and color at a given position, and to sort a list of candidates in descending order according to the ETSD score. An example of the ECMG for the aforementioned movie poll can be seen in Figure 1.C.

5 Experimental Evaluation

So as to perform the evaluation we have conducted an experiment comparing the aforementioned three visualization techniques. During the experiment different datasets were shown to the users using our visualization techniques. The users were then asked to make observations about the information they could extract. These observations have then been quantified using an insight-based method [18, 17]. We have used this method because a classical task-based evaluation methodology did not comply with the nature of the voting problem, and thus would have not properly evaluated the relevance of our technique.

5.1 Data Sets and Visualizations

In order to evaluate our technique we used data from three different voting polls. The data sets are the preference profiles for specific elections. The first data set was about a movie poll, where we asked people (voters) to rank their favorite movies (candidates) from the 70s and later. The preference profile can be seen in the appendix (Figure 2). The value (number) in the entry $a_{i,j}$ of the matrix gives the ranking of candidate j in the preference of voter i, *e.g.*, voter "#2" ranks "The Shawshank Redemption" as her 8$^{\text{th}}$ choice while "Forrest Gump" as her first choice.

The two other data sets are synthetic preference profiles: we chose a number of m candidates, assigned to them a probability distribution of being chosen, and then built N votes by sampling the candidates without replacement according to the probability distribution. This gives us random preferences, but with a bias that makes some candidates most likely to be at the beginning of the preferences, and thus most likely to be a Condorcet winner.

The results of the elections for the aforementioned polls were displayed using the proposed techniques, *i.e*, the Pairwise Comparison Matrix (Section 4.1), the Weighted Majority Graph (Section 4.2) and the Edge-Compressed Majority Graph (Section 4.3).

5.2 Participants and Protocol

In order to evaluate the visualization techniques we designed the following protocol for our experiment. It was conducted with six participants. None of them had any background in Social Choice nor in Voting Theory. The protocol consisted of three phases. In the first phase, a short tutorial of 15 minutes was given to the participants explaining all the necessary information in order for them to get familiar with the experimental setting. We started off with the analysis of the key aspects of Voting Theory, *i.e.*, the voting problem and the Condorcet criterion, and then we presented a brief introduction of the three visualization techniques. In the second phase of our experiment, that lasted about 30 minutes, each user was given a specific poll from the data sets described above. The results of the poll were shown to the user with each of the three techniques one at a time in a random order. When the synthetic data sets were used, we changed the labels of the candidates when the new technique was shown so that the participants would not be biased about the data set that they had just previously seen. The users were instructed to look carefully at the techniques and report their observations (insights) with all the information they could extract out of the polls. When the users felt they could not gain any other information out of one technique we proceeded to the next one. We took notes during the session but also videotaped the whole procedure for later identification and analysis of all individual occurrences of insights. In the third and final phase of the experiment, the users were asked to make general comments about the techniques and to express their personal preference regarding the technique they found the most appropriate one for the data sets.

5.3 Results

In order to analyze the results of our experiment we have studied the insights gathered by the participants. We present our findings about the insight characteristics and the participants' general comments.

5.3.1 Evaluation on Insight Metrics

During our analysis we used various insight metrics which are summarized in Table 1 and are thoroughly analyzed in the following paragraphs.

	PCM	WMG	ECMG
Insights (total)	304	218	315
Insight Categories			
Ranking	35	37	39
Comparisons	203	131	302
Margins	42	38	
Hypotheses/Remarks	7	5	16
Inaccurate	3	4	0
Domain Value	950	604	1310
Avg. time of 1st insight (sec)	13	10	5
Total time	6 : 51	6 : 43	5 : 58
Participant's evaluation (out of 10)	7.16	4.83	8.83

Table 1: Overview of the Results.

Insights We measured the total number of insights, *i.e.*, the total number of observations about the information a user could extract. We used the following 4 categories (ranking, margin, comparisons, hypotheses) in order to categorize the insights and assigned specific domain value points for each category reflecting the importance of the information extracted. The insights regarding information about the ranking of the candidates, as well as the insights about the difference in the margin of victory/loss in terms of voters for a given pairwise comparison were given 1-2 points. The insights regarding information about the results of pairwise comparisons between candidates were given 3-4 points. More specifically, we gave 3 points if a single pairwise comparison was detected while 4 points were given for each pairwise comparison if multiple comparisons were spotted at a glance. Last, the insights regarding hypotheses and more critical remarks were given 5 points. For example, the most observed critical remark by the participants was the identification of Condorcet Paradox. Recall that the participants didn't have any background on voting theory but identified the cycle(s) between the candidates and speculated that there is no winner in this(ese) case. In conclusion, we noticed that the total number of insights is slightly greater in the ECMG compared to the PCM and both of these techniques have a huge difference with the WMG.

Categories of Insights Given that the number of insights is fairly equal for the ECMG and the PCM techniques, a clearer conclusion can be made by looking at the insights per category. We can see that the number of insights for the ranking of candidates is approximately the same for all the techniques, as well as the number of insights regarding the margins. We see that the ECMG has a strong lead when compared to the other two techniques in the number of insights concerning the information about the pairwise comparisons. We can see that ECMG has 302 insights while PCM has 203 and WMG has 131. A similar pattern is also seen if we categorize the observations of pairwise comparisons into single or not. The 276 out of 302 comparisons observed with the ECMG were about multiple comparisons spotted at a glance while PCM had 189 and WMG 110. Regarding the number of hypotheses and the critical remarks, we notice that ECMG is much superior compared to the two other techniques (16 versus 7 and 5). Given the importance of the pairwise comparisons in the Voting Theory and the qualitative nature of the insight method it is clear that the ECMG technique outperforms the other two techniques when taking into account the categorization of the insights.

Domain Value The domain value is critical in evaluating the techniques as it reflects (in a quantitative way) the quality of the observations. We use the aforementioned category-based point system for measuring the quality of each insight, and sum up all the points in order to compute the total domain value. This quantification of the information extracted by the participants allows to take into account the significance of the pairwise comparisons and other principal elements of the Voting Theory. The ECMG technique achieved the highest score among all three techniques, *i.e.*, 1310 points. The PCM comes second with a score of 950 points and the WMG technique performs the lowest score, reaching a total of 604 points. We can remark here that the differences between the scores of the techniques are significant and thus ECMG outperforms the other two techniques in terms of the total domain value.

Average Time of First Insight This metric measures the time that passes from the beginning of the session until the occurrence of the first insight made by the user. A shortest time on observing the first insight implies that participants are able to grab information quicker and thus have a shorter learning time. In 85% of the cases, the first insight that was made by the users was about the candidate who is the winner of the election. The participants needed half the time to observe the first insight when the ECMG technique was used in comparison with the WMG and almost a third of the time compared to the PCM. Given that a high percentage of the first insight is referring to a crucial information about the elections, we can imply that the ECMG technique is dominant over the other techniques regarding the time needed for the first data extraction.

Total Time and Insights per Minute When measuring time it is also important to measure the total time and the insights observed per minute. The total time measures the time needed for the participants to feel that they could not gain any further information out of the graph. Having the total time lower for one technique has two interpretations. The first is that the technique is more efficient because users needed less time to complete the extraction of information. The second is that users gave up because of lack of interest as they felt they could not gain any more information. The results reveal that the ECMG technique was faster and also produced more insights giving an average 52.8 insights per minute, while PCM had 44.4 and WMG had 32.5. After analyzing the results we conclude that ECMG is the most efficient technique from time perspective because time was less and despite that, more insights were found giving the technique the highest ratio of insights/minute.

5.3.2 Participant Comments on Visualization Techniques

In order to complete the evaluation of the visualization techniques we asked users to comment and express their opinion (positive or negative) according to the difficulties they encountered. At the end, the participants were asked to perform an evaluation on the techniques by ranking them and assigning a score according to their level of satisfaction. The findings for each technique are the following.

Pairwise Comparison Matrix The PCM technique achieved a mean rating of 7.16/10 according to the preferences of the users, which ranks the technique second. The technique was the most preferable one according to one user, *i.e.*, 16.7% of the users. A general comment made by most of the participants was that it is a good technique to show the results of single pairwise comparisons and the margins of difference but it is not good to see the Condorcet cycles. Most users said that it is a good way to visualize ranking. Generally most participants agreed that you can extract the same amount of information as in ECMG but you need longer time. There were some rare comments about the color degradation which caused difficulty in identifying the results of the pairwise comparisons when the margins of differences between the candidates were narrow. Lastly, one user did not find this technique as visually attractive as the other ones.

Weighted Majority Graph The WMG technique was the least preferable one as it was ranked third by all participants and scored a rating of 4.83/10. Most of the participants (83%) noticed that the presentation is unclear. They commented that it was complex to see the edges as they were crossing each other and it was also hard to distinguish the arrows on the candidates. Therefore they said that they needed time and concentration to figure out the pairwise comparisons. Another major criticism was about the ranking according to the extended Tideman's simplified Dodgson rule, which they found difficult to see.

Edge-Compressed Majority Graph The ECMG technique was the best one according to all participants but one and acquired a mean score of 8.83/10. Regarding the ranking of the candidates according to extended Tideman's simplified Dodgson rule most users (83%) agreed that it was easy to detect, apart from one user who found it difficult to distinguish the ties in ranking. The great majority of the users (83%) found the technique as a natural and straightforward method for displaying the results of the pairwise comparisons between the candidates. In particular, they said that it is really clear to follow the fewer edges and the "up to down" layout of the graph. Also half of the users added that it was a nice technique for representing the voting cycles and the Condorcet paradox.

5.3.3 Discussion of Results

Overall, the results showed the superiority of the ECMG over the other techniques. The ECMG technique produced a slightly higher number of insights compared to the PCM but the insights were of higher value because they were mostly referring to pairwise comparisons and new hypotheses. ECMG had a mean domain value of 4.16 per insight while PCM had 3.13 and WMG had 2.77. This finding was also confirmed by the participants' comments about the visualization of comparisons in all three techniques.

We found very promising that the hypotheses made by the users confirmed our initial expectations that the ECMG can provide an alternative method for interpreting the ranking, given the various discrepancies in voting rules. The experimental results revealed that the participants had the tendency to give their own rankings which sometimes were different from the extended Tideman's simplified Dodgson's rule ranking. Therefore, we think that, from a social point of view, someone can use our technique to obtain the ranking according to her assumptions and priorities given that voting rules can produce various rankings for the same poll. In this way we can confirm the expectation imposed at the beginning of our research that the ECMG is the intermediate step between the individual preferences and the collective ranking.

Furthermore ECMG was the only technique without any incorrect insights. It also had more than twice as many critical remarks and hypotheses than both PCM and WMG. Concerning time, the ECMG surpassed the other two techniques because it required less time for the first insight to be observed, and also less time in total while more insights were found. The time perspective is also widely acknowledged by the participants, who found ECMG the fastest way of observing the results. Finally, the participants' personal evaluation and ranking of all three techniques revealed the dominance of ECMG.

6 Conclusion and Future Work

In this paper, we have proposed three techniques for graphically representing the weighted majority graph of an election. Two of them, PCM and WMG, are classical visualization techniques for directed graphs. The last one, ECMG is original and based on edge-compression techniques. Experiments carried out with non specialist users show that the ECMG yields a lot more insights than the two other techniques, and hence that it is more efficient in providing information about the structure of the collective preference.

As future work, we plan to integrate the ECMG to the set of visualization proposed on the Whale[3] web application [5], so as to make it directly available to the users of this system. We also plan to investigate visualization techniques for other families of voting rules that are not based on pairwise comparisons. Another important challenge is to find interesting techniques that represent basic theorems and notions of Social Choice *e.g.*, visualizing the discrepancies that occur for the same preference profile among different voting rules.

References

[1] J. Banks. Sophisticated voting outcomes and agenda control. *Social Choice and Welfare*, 1(4):295–306, 1985.

[2] J. Bartholdi, C. A. Tovey, and M. A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6:157–165, 1989.

[3] N. Betzler, R. Bredereck, and R. Niedermeier. Theoretical and empirical evaluation of data reduction for exact Kemeny Rank Aggregation. *Autonomous Agents and Multi-Agent Systems*, 28(5):721–748, 2014.

[4] D. Black. *Theory of Committees and Elections.* Cambridge University Press, 1958.

[5] S. Bouveret. Whale^3. `http://whale3.noiraudes.net/`.

[6] C. Brewer. Colorbrewer 2.0. `http://colorbrewer2.org/`.

[7] I. Caragiannis, C. Kaklamanis, N. Karanikolas, and A. D. Procaccia. Socially desirable approximations for Dodgson's voting rule. *ACM Transactions on Algorithms*, Volume 10(Issue 2), February 2014.

[8] A. Copeland. A reasonable social welfare function. In *Seminar on applications of mathematics to social sciences, University of Michigan*, 1951.

[9] N. de Condorcet. *Essai sur l'application de l'analyse à la probabilité de décisions rendues à la pluralité de voix.* Imprimerie Royal, facsimile published in 1972 by Chelsea Publishing Company, New York. edition, 1785.

[10] T. Dwyer, N. H. Riche, K. Marriott, and C. Mears. Edge compression techniques for visualization of dense directed graphs. *IEEE Trans. Vis. Comput. Graph.*, 19(12):2596–2605, 2013.

[11] Electowiki. Maximal elements algorithms. `http://wiki.electorama.com/wiki/Maximal_elements_algorithms`, 2005.

[12] P. Fishburn. Condorcet social choice functions. *SIAM Journal of Applied Mathematics*, 33:469–489, 1977.

[13] I. L. Grivko and V. Levchenkov. Intrinsic properties of the self-consistent choice rule. *Automation and Remote Control*, 55:689–697, 1994.

[14] J.-F. Laslier. *Tournament Solutions and Majority Voting.* Springer-Verlag Berlin Heidelberg, 1997.

[15] V. Levchenkov. Social choice theory : a new sight. *Preprint of the Institute for System Analysis, Moscou.*, 1992.

[16] N. Miller. A new solution set for tournaments and majority voting: Further graph-theoretical approaches to the theory of voting. *American Journal of Political Science*, 24(1):68–96, 1980.

[17] C. North, P. Saraiya, and K. Duca. A comparison of benchmark task and insight evaluation methods for information visualization. *Information Visualization*, 10(3):162–181, 2011.

[18] P. Saraiya, C. North, and K. Duca. An insight-based methodology for evaluating bioinformatics visualizations. *IEEE Trans. Vis. Comput. Graph.*, 11(4):443–456, 2005.

[19] J. H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6):1027–1041, 1973.

[20] N. Tideman. *Collective Decisions and Voting.* Ashgate, 2006.

[21] H. P. Young. Extending Condorcet's rule. *Journal of Economic Theory*, 16:335–353, 1977.

Nikos Karanikolas
LIG, Université Grenoble Alpes
Grenoble, France
Email: Nikolaos.Karanikolas@imag.fr

Renaud Blanch
LIG, Université Grenoble Alpes
Grenoble, France
Email: Renaud.Blanch@imag.fr

Sylvain Bouveret
LIG, Université Grenoble Alpes
Grenoble, France
Email: Sylvain.Bouveret@imag.fr

Appendix

The preference profile for the movie poll The data set used in the following figures is the movie poll data set described in Section 5.1 and depicted in the following figure 2. In this election setting we have 22 voters with each one having a numerical id and 10 candidates corresponding to the movies.

	The Shawshank Redemption	The Godfather	The Dark Knight	Pulp Fiction	Schindler's List	The Lord of the Rings	Fight Club	Star Wars	Forrest Gump	The Matrix
#1	4	9	3	6	7	8	1	10	5	2
#2	8	5	10	3	6	7	2	9	1	4
#3	6	9	7	4	10	5	1	8	2	3
#4	10	9	3	7	8	5	1	4	2	6
#5	9	8	2	10	4	1	6	3	7	5
#6	8	9	3	2	10	5	1	4	7	6
#7	10	9	8	5	6	7	1	2	3	4
#8	10	6	4	1	9	7	2	5	8	3
#9	10	6	5	1	3	9	7	2	8	4
#10	3	6	10	4	2	8	1	9	5	7
#11	10	7	9	5	4	6	8	1	3	2
#12	10	9	8	7	1	5	3	6	2	4
#13	10	6	9	7	3	1	5	4	8	2
#14	10	9	4	7	8	6	1	5	2	3
#15	9	3	8	1	10	7	4	2	5	6
#16	10	8	3	7	4	6	1	9	5	2
#17	10	3	5	2	4	9	8	1	7	6
#18	2	3	6	10	1	7	9	8	4	5
#19	1	4	7	2	3	8	10	9	5	6
#20	10	9	8	4	7	1	3	2	6	5
#21	4	2	10	1	9	6	3	5	8	7
#22	8	5	9	4	1	3	10	7	2	6

Figure 2: Preference Profile for the Movie poll.

Algorithms The pseudocode below for Algorithm 1 describes our main technique used for the computation of the ECMG visualization. Algorithm 2 is the main drawing procedure and is called by Algorithm 1 in order to draw the Majority Graphs computed in each step of the while loop.

Algorithm 1 Edge-Compressed Majority Graph Visualization Technique

procedure Circular Smith Set
  $C \leftarrow$ the set of candidates
  $i \leftarrow 0$
  while $C \neq \emptyset$ **do**
    $S_i \leftarrow$ ComputeSmithSet(C)
    if the graphical layout is empty **then** ▷ Algorithm 2
      draw the $MG(S_i)$ on the top of the graphical layout
    else
      draw the $MG(S_i)$ on the graphical layout under $MG(S_{i-1})$
    end if
    Let $C = C \setminus S_i$
    $i \leftarrow i + 1$
  end while
end procedure

Algorithm 2 Majority graph visualization

```
procedure DrawMG($S$)
    **Node's attributes:**                                  ▷ Depend on candidate's ranking
    $c \leftarrow$ candidate
    $D \leftarrow$ node's diameter
    $color \leftarrow$ node's color
    $position_{x,y} \leftarrow$ node's position in x,y axis
    if size $(S) = 1$ then                                  ▷ Smith Set has a single member
        DrawNode($c$,$D$,$color$,$position_x$,$position_y$)
    else                                ▷ Smith Set consisting of a cluster of candidates
        $offset \leftarrow$ spatial distance between 2 consecutive nodes $S$
        $N \leftarrow$ number of candidates in $S$
        $Sorted_i \leftarrow$ SortNodes($VotingRule$)
        if nodes $\in S$ have linear (total) order then
            $w, h \leftarrow$ rectangle's dimensions
            Draw a rectangle with dimensions $w, h$
            for $i \leftarrow 1, N$ do                      ▷ Draws $S$ nodes inside rectangle
                if $score(Sorted_i = Sorted_{i-1})$ then    ▷ 2 or more candidates have the same score
                    $position_x \leftarrow position_x + offset$
                    DrawNode($Sorted_i$,$D$,$color$,$position_{x,y}$)   ▷ nodes are drawn parallel
                else
                    $position_y \leftarrow position_y + offset$
                    DrawNode($Sorted_i$,$D$,$color$,$position_{x,y}$)▷ the node with higher score is drawn lower
                end if
            end for
        else                                        ▷ voting cycles between candidates occur
            $D \leftarrow$ Circle's diameter
            Draw a circle with diameter $D$
            for $i \leftarrow 1, N$ do                      ▷ Draws $S$ nodes inside circle
                $position_{x,y} \leftarrow position_{x,y} + offset$
                DrawNode($Sorted_i$,$D$,$color$,$position_{x,y}$)   ▷ when $score_i \geq score_{i-1}$, $i$ is drawn lower
            end for
        end if
        Draw the edges between the nodes of $S$
    end if
end procedure
```