

Explaining Tournament Solutions with Minimal Supports

Clément Contet^{1,2}, Umberto Grandi^{1,3}, Jérôme Mengin^{1,2}

¹Institut de Recherche en Informatique de Toulouse (IRIT)

²Univeristé de Toulouse

³Université Toulouse Capitole

{clement.contet, umberto.grandi, jerome.mengin}@irit.fr

Abstract

Tournaments are widely used models to represent pairwise dominance between candidates, alternatives, or teams. We study the problem of providing certified explanations for why a candidate appears among the winners under various tournament rules. To this end, we identify minimal supports—minimal sub-tournaments in which the candidate is guaranteed to win regardless of how the rest of the tournament is completed (that is, the candidate is a necessary winner of the sub-tournament). This notion corresponds to an abductive explanation for the question, “Why does the winner win the tournament?”—a central concept in formal explainable AI. We focus on common tournament solutions: the top cycle, the uncovered set, the Copeland rule, the Borda rule, the maximin rule, and the weighted uncovered set. For each rule we determine the size of the smallest minimal supports, and we present polynomial-time algorithms to compute them for all solutions except for the weighted uncovered set, for which the problem is NP-complete. Finally, we show how minimal supports can serve to produce compact, certified, and intuitive explanations for tournament solutions.

1 Introduction

Tournaments are well-known mathematical structures to represent pairwise contests between alternatives, be they a set of candidates in an election, different teams in a sport league, or competing political proposals. A number of tournament solutions have been proposed and studied in the literature to aggregate the information represented in such graphs and select a set of winning alternatives. For an introduction, see Brandt, Brill, and Harrenstein (2016) and Fischer, Hudry, and Niedermeier (2016).

As simple tools for collective decision making, tournament solutions are valuable for digital democracy. However, their adoption will remain limited to small-scale, low-stakes elections until further efforts are made to make them more transparent and accessible to users (Grossi et al. 2024). Leveraging techniques from explainable AI (XAI), in this paper we propose to supplement tournament solutions with a compact certified explanation, enabling voters to efficiently verify the winning alternative while at the same time gaining a clearer understanding of the tournament solution.

XAI aims at helping human users to comprehend and trust the outputs of AI systems. To this end, multiple properties have been introduced to define effective explanations (see, e.g., Miller 2019, for an introduction). First, an explanation should contain the right quantity of information, ensuring sufficiency without excess (see, e.g., the irreducibility and validity properties by Amgoud and Ben-Naim 2022, and Amgoud, Cooper, and Salim 2024). Second, to foster trust, only true statements supported by evidences should be used (Kulesza et al. 2013). Third, the information provided must be relevant to the problem at hand. Finally, an explanation should be brief, organized and non ambiguous (see, e.g., empirical evidence by Delaunay et al. 2025). Grice (1975) refers to these four criteria as *quantity*, *quality*, *relation* and *manner*.

Since tournament solutions are white-box models, a natural candidate for an explanation is to provide voters with the full information on the vote counts, to be able to recompute the solution and verify it. However, this approach requires more communication effort than necessary, does not provide any new insight to understand the outcome, and is only accessible to experts or at least computer-savvy people for the more advanced tournament solutions.

In the literature, three main options have been proposed to explain and justify the choice returned by a given tournament solution. First, in social choice and welfare economics (see, e.g., Arrow, Sen, and Suzumura 2010) one is presented with a set of desirable axiomatic properties or an intuitive mathematical structure that constitutes the normative justification of a tournament solution (Cailloux and Endriss 2016; Boixel, Endriss, and de Haan 2022; Nardi, Boixel, and Endriss 2022; Peters et al. 2020). Yet, explaining why a decision is good does not necessarily help in explaining the actual decision-making process, since multiple outcomes can be supported by different good reasons. Alternatively, voters can be presented with detailed or aggregated data on their expressed preferences, in the form of statements or features of the election, manually computed or automatically extracted (Suryanarayana, Sarne, and Kraus 2022). Although more accessible, the explanations produced in this way offer no logical guarantee and may be insufficient to justify a specific result. Finally, a recent approach by Contet, Grandi, and Mengin (2024) experiments with the use of abductive and contrastive explanations from formal explanations in ma-

Tournament Solution	Computing an SMS	Lower bound on SMS size	Upper bound on SMS size
Top Cycle	$\mathcal{O}(m^2)$	$m - 1$	$m - 1$
Uncovered Set	$\mathcal{O}(m^2)$	$m - 1$	$m - 1$
Copeland	$\mathcal{O}(m^2)$	$m - 1$	$(m - 1) \lfloor \frac{m-1}{2} \rfloor$
Borda	$\mathcal{O}(m^2 \log n)$	$\lceil n \frac{(m-1)^2}{m} \rceil$	$(m - 1) \lfloor n \frac{m-1}{2} \rfloor$
Maximin	$\mathcal{O}(m^2 \log n)$	$\lceil \frac{n}{2} \rceil (m - 1)$	$n(m - 1)$
Weighted Uncovered Set	NP-complete	$n + m - 2$	$(n + 1)(m - 1)$

Table 1: Overview of our results for various tournament solutions with n voters and m candidates. All bounds are tight. Closed-form expressions of the SMSs size are also provided for every solution but the weighted uncovered set.

chine learning (see, e.g., Shih, Choi, and Darwiche 2018, Darwiche and Hirth 2020, Ignatiev et al. 2020, Marques-Silva 2022), introducing them to a matrix-based voting setting. However, these formal explanations often are not user-friendly in view of their size and lack of visual structure.

In this paper we build on the latter proposal to introduce the concept of *minimal supports (MS)* for tournaments solutions, which encapsulate the core reasons why a tournament winner belongs to the winner set. We define minimal supports of a winning alternative w for a given tournament solution S on a tournament G as the set of all minimal sub-tournaments of G for which w is a necessary winner, i.e., w is the winner of S in all completions of the MS. MSs provide a rigorous base to explain the outcome of a tournament, as they correspond to *abductive (or prime implicant) explanations* to the question “Why is w the S -winner of tournament G ?”.¹ To show how our explanations can be made accessible, we provide efficient algorithms to select the *smallest MSs (SMS)* and we show how to build rigorous, compact and accessible explanations in natural language using MSs.

After introducing preliminary notions in Section 2, we introduce in Section 3 the concept of *minimal support* for tournaments solutions. While the problem of finding smallest abductive explanations is often intractable (without constraints on the classifier, the decision problem is Σ_2^P -hard (Liberatore 2005)), in Section 4 we provide polynomial algorithms for computing *smallest minimal supports* and closed-form expressions of their size for the top cycle (TC), the uncovered set (UC), the Copeland rule (CO), the Borda rule (BO), and the maximin rule (MM). We also show that the problem is NP-complete for the weighted uncovered set (wUC). Additionally, we bound the size of the smallest MSs and show that have underlying structure. In Table 1 we provide a detailed overview of our technical results. Finally, in Section 5 we show how MSs can be used to deliver certified, compact, and intuitive explanations in natural language.

¹Abduction has been defined in a variety of ways in XAI for ML and in logic-based explanations. In particular, sometimes it is limited to “why not” questions (see, e.g., Eiter and Gottlob 1995). In this paper, we follow the terminology of Ignatiev et al. (2020).

1.1 Related Work


We identified three approaches in the literature aiming to explain voting outcomes. First, in the realm of voting with rankings, Suryanarayana, Sarne, and Kraus (2022) compared experimentally crowdsourced arguments with algorithmic explanations using features of a preference profile such as various notions of score. They showed that both approaches produce comparable results in terms of acceptance and legitimacy of the process, and identified a positive effect on the acceptance of the winner by the less satisfied users. Second, a stream of papers mentioned earlier (Cailloux and Endriss 2016; Boixel, Endriss, and de Haan 2022; Nardi, Boixel, and Endriss 2022; Peters et al. 2020) developed a logical calculus based on axiomatic properties to justify why a candidate should be the winner in a preference profile, building on the normative justification of voting rules. Third, building on recent work in formal XAI, Contet, Grandi, and Mengin (2024) explored the use of abductive explanations in voting with rankings for the Borda rule.

In formal explainability, abductive explanations have a dual structure called contrastive explanations (Ignatiev et al. 2020). If abductive explanations answer to questions such as “Why is a the result of classifier F on input I ?”, contrastive explanations answer questions of the form “Why is b not the result of classifier F on input I ?”. These counterfactual questions are answered by identifying minimal sets of input features that have the potential of reverting the classifier result, or that, if left unchanged, can guarantee that the same outcome is reproduced. When applied to voting, contrastive explanations are strongly related to the problem of *bribery*, studied under various settings ranging from entire changes of a voter’s ballot to limited modifications (see, e.g., Faliszewski, Hemaspaandra, and Hemaspaandra 2009, and Elkind, Faliszewski, and Slinko 2009). The setting considered in this paper, where each pairwise comparison can be individually changed ignoring the transitivity of the preferences, was introduced by Faliszewski et al. (2009).

The study of bribery in tournaments has lead Brill, Schmidt-Kraepelin, and Suksompong (2020) to define the notion of *margin of victory*, later extended by Döring and Peters (2023), in an effort to refine the outcome of tournament solutions. Since contrastive explanations are answers to “why not?” questions, they are great tools to show why a losing alternative was not selected. However, explaining individually why each losing alternative did not win is not necessarily a good explanation of why the winner was elected (Lipton 1990).

While bribery is well-studied, its dual problem—defending against it, akin to finding abductive explanations—could have been explored in similar depth. However, to our knowledge, only Chen et al. (2021) studies the computational complexity of protection against both constructive and destructive bribery, showing that protection is generally harder than the bribery problem itself.

Our approach is in line with the work of Beynier et al. (2024), who provide explanations for the lack of local envy-free allocations in fair division using formal methods akin to abductive explanations and present them with graphs.



(a) (unweighted) tournament G (b) 5-weighted tournament G_w

Figure 1: A tournament and a 5-weighted tournament.

2 Preliminaries

Tournaments are well-studied structures used to represent compactly pairwise comparisons among alternatives. In the following presentation we use the framework of Aziz et al. (2015). Formally, a *partial tournament* is a pair $G = (\mathcal{C}, E)$ where \mathcal{C} is a nonempty finite set of candidates or alternatives and $E \subseteq \mathcal{C} \times \mathcal{C}$ is an asymmetric relation on \mathcal{C} , i.e., $(y, x) \notin E$ whenever $(x, y) \in E$. A *tournament* G is a partial tournament (\mathcal{C}, E) for which E is also complete, i.e., either $(x, y) \in E$ or $(y, x) \in E$ for all distinct $x, y \in \mathcal{C}$.

Let the *number of voters* n be a strictly positive integer. A *partial n -weighted tournament* is a pair $G = (\mathcal{C}, \mu)$ where \mathcal{C} is a nonempty finite set of candidates and $\mu : \mathcal{C} \times \mathcal{C} \rightarrow \{0, \dots, n\}$ a weight function such that for all distinct $x, y \in \mathcal{C}$, $\mu(x, y) + \mu(y, x) \leq n$. We also assume $\mu(x, x) = 0$ for every candidate x . A (*complete*) *n -weighted tournament* satisfies for all distinct $x, y \in \mathcal{C}$, $\mu(x, y) + \mu(y, x) = n$. Naturally, (unweighted) tournaments can be seen as 1-weighted tournaments.

Example 1. Consider tournament $G = (\mathcal{C}, E)$ in Figure 1a, the arrow from node a to node b indicates that a is preferred to b , or, that a beats b in a pairwise comparison. In the 5-weighted tournament $G_w = (\mathcal{C}, E_w)$ in Figure 1b, the “3” on the arrow from a to b means that 3 voters prefer a to b , or, that a beats b in 3 pairwise comparisons.

In the remainder of this paper, edges of weighted tournament with a weight of zero are not represented.

A *tournament solution* (or rule) is a function S mapping a tournament $G = (\mathcal{C}, E)$ (or a weighted tournament $G = (\mathcal{C}, \mu)$) to a non-empty set of candidates $S(G) \subseteq \mathcal{C}$. In this paper we focus on six common and widely studied rules.

Let $G = (\mathcal{C}, E)$ be a complete tournament. A nonempty subset of candidates $A \subseteq \mathcal{C}$ is called *dominant* in G if for every candidate $x \in A$, for all candidates $y \in \mathcal{C} \setminus A$, $(x, y) \in E$. A candidate $x \in \mathcal{C}$ is said to *cover* another candidate $y \in \mathcal{C}$ if for all $z \in \mathcal{C}$, $(y, z) \in E \implies (x, z) \in E$.

- The *top cycle* (TC) is the (unique) minimal dominant nonempty subset of candidates of G , i.e., the top strongly connected component of G . Alternatively, it can also be defined as the set of candidates that can reach every other candidate via a directed path in G .
- The *uncovered set* (UC) is the (unique) nonempty subset of candidates not covered by any other candidate in G . Alternatively, it can also be defined as the set of candidates that can reach every other candidate via a directed path in G of length at most two.

- The *Copeland score* of a candidate $c \in \mathcal{C}$ in $G = (\mathcal{C}, E)$ is $\sigma_{\text{CO}}(c, G) = |\{c' : c' \in \mathcal{C}, (c, c') \in E\}|$. The *Copeland rule* (CO) selects candidates that have a maximal Copeland score.

Given n voters, let $G = (\mathcal{C}, \mu)$ be a complete n -weighted tournament. A candidate $x \in \mathcal{C}$ is said to *weighted cover* another candidate $y \in \mathcal{C}$ if for all $z \in \mathcal{C}$, $\mu(x, z) \geq \mu(y, z)$.

- The *Borda score* of a candidate $c \in \mathcal{C}$ in $G = (\mathcal{C}, \mu)$ is $\sigma_{\text{BO}}(c, G) = \sum_{c' \in \mathcal{C}} \mu(c, c')$. The *Borda rule* (BO) selects candidates that have a maximal Borda score.
- The *maximin score* of a candidate $c \in \mathcal{C}$ in $G = (\mathcal{C}, \mu)$ is $\sigma_{\text{MM}}(c, G) = \min_{c' \in \mathcal{C} \setminus \{c\}} \mu(c, c')$. The *maximin rule* (MM) selects candidates that have a maximal maximin score.
- The *weighted uncovered set* (wUC) is the nonempty subset of candidates that are not weighted covered by any other candidate in G . A path-based alternative definition also exists (see Lemma 7).

Observe that these six tournament solutions can be organized in two families. First, there are the *path-based solutions*: TC, UC and wUC which are solutions associated to some notion of paths. Second, CO, BO and MM are rules which score candidates where each candidate’s score can be computed using only the pairwise comparisons between it and the other candidates. We call them *myopic-score solutions*. This distinction will become useful in Section 3.3 when discriminating between MSs.

3 Minimal Supports for Tournaments

This section adapts the definitions from the literature on formal explanations to the case of tournaments, and proposes three principles to choose among minimal supports.

3.1 MSs for Unweighted Tournaments

Abductive reasoning applied to tournament solutions explains the winner of a tournament by exhibiting a minimal subset of the tournament such that the winner remains unchanged. Let $G = (\mathcal{C}, E)$ and $G' = (\mathcal{C}, E')$ be two partial tournaments, we say that G' is an *extension* of G , denoted $G \subseteq G'$, if $E \subseteq E'$. If E' is complete, G' is called a *completion* of G , and we write $[G]$ for the set of completions of G . To obtain a formal definition we use the concept of necessary winner (Konczak and Lang 2005).

Definition 1. Given a partial tournament G , a tournament solution S and a candidate $c \in \mathcal{C}$, c is a necessary winner of S in G if for every completion $G' \in [G]$ we have that $c \in S(G')$. We write $c \in \text{NW}_S(G)$.

We now define minimal supports for tournaments:

Definition 2. Given a tournament $G = (\mathcal{C}, E)$, a tournament solution S and a winning candidate $w \in S(G)$, a minimal support (MS) for $w \in S(G)$ is a partial tournament $G' \subseteq G$ such that:

- $w \in \text{NW}_S(G')$
- G' is \subseteq -minimal, i.e., all partial tournaments $G'' \subsetneq G'$ are such that $w \notin \text{NW}_S(G'')$.




Figure 2: MSs for $a \in UC(G)$ from Figure 1a.




Figure 3: MSs for $a \in MM(G_w)$ from Figure 1b.

Example 2. Consider tournament G in Figure 1a. The partial tournament \mathcal{X} represented in Figure 2a is an MS for $a \in UC(G)$. Note that $\mathcal{X} \subseteq G$. What \mathcal{X} shows is that a is not covered by b since a is preferred to b ; a is not covered by c since a is preferred to b and b is preferred to c ; a is not covered by d since a is preferred to d . Hence, a is in the uncovered set. Similar reasoning holds for \mathcal{Y} in Figure 2b.

3.2 MSs for Weighted Tournaments

The case of weighted tournaments can be treated analogously by simply defining an inclusion relation for weighted tournaments. Let $G = (\mathcal{C}, \mu)$ and $G' = (\mathcal{C}, \mu')$ be two partial n -weighted tournaments, we say that G' is an *extension* of G , denoted $G \subseteq G'$, if $\mu(x, y) \leq \mu'(x, y)$ for all $(x, y) \in \mathcal{C} \times \mathcal{C}$. A *completion* of G is a complete n -weighted extension, and both Definitions 1 and 2 apply.

Example 3. Consider G_w in Figure 1b, both \mathcal{I} and \mathcal{J} , represented in Figure 3, are MSs for $a \in MM(G_w)$. Observe that $\mathcal{I} \subseteq G_w$ and $\mathcal{J} \subseteq G_w$. \mathcal{I} shows that a achieves a maximin score of 3 and that no other candidate can beat it since they all end up with 3 defeats against a . \mathcal{J} only secures a maximin score of 2 for a and uses additional pairwise comparisons among other candidates to show that they cannot outperform a .

3.3 Selecting among MSs for Explanations

For a given tournament there might be an exponential number of different minimal supports since MSs are an instance of a Sperner family, a family of subsets which do not contain each others (Sperner 1928). Therefore in this section we introduce criteria to select the most promising MSs from an explainability point of view.

Every MS can be a certificate for an explanation since by definition it is subset minimal and presents precisely enough

information to guarantee the outcome of the tournament. Hence, looking back at the Gricean principles (Grice 1975) mentioned in the introduction, all MSs are optimal for quality. However, we can still use the manner criteria to discriminate between MSs since some can be smaller than others, i.e., rely on a smaller total amount of information. This is interesting from the perspective of formal verification since it is easier to verify the correctness of smaller certificates (imagine you have to call back participants or re-watch the games between contestants to verify that the registered preferences are correct).

To measure this amount of information in a general context, we take inspiration from the bribery setting and we fix the smallest atom of information to be the individual swap, i.e., the preference of one voter in the pairwise comparison between two candidates.

Definition 3. Given n voters, a partial n -weighted tournament $G = (\mathcal{C}, \mu)$, and a winning candidate $w \in S(G)$, we define the size of an MS $\mathcal{X} = (\mathcal{C}, \mu_{\mathcal{X}})$ for $w \in S(G)$ as $|\mathcal{X}| = \sum_{(c, c') \in \mathcal{C}^2} \mu_{\mathcal{X}}(c, c')$. \mathcal{X} is a smallest minimal support (SMS) for $w \in S(G)$ if and only if for all MSs \mathcal{Y} for $w \in S(G)$, we have $|\mathcal{X}| \leq |\mathcal{Y}|$.

Example 4. In Figure 2, $|\mathcal{X}| = 3$, $|\mathcal{Y}| = 3$ and both are SMSs for $a \in UC(G)$. Both \mathcal{I} and \mathcal{J} in Figure 3 are MSs for $a \in MM(G_w)$. However, \mathcal{I} is smaller than \mathcal{J} as $|\mathcal{I}| = 9$ and $|\mathcal{J}| = 13$. Only \mathcal{I} is an SMS for $a \in MM(G_w)$.

Smallest MSs for weighted tournaments can be viewed as MSs such that no strictly smaller MS exists for the ℓ_1 -norm ($\|x\|_1 = \sum_{i \in I} |x_i|$) of the tournament weight vector. For unweighted tournaments this reduces to minimizing the number of edges.

To further select MSs, we propose two additional criteria to choose among SMSs depending of the type of tournament solution. For the path-based solutions, it is natural for simplicity purposes to look for SMSs containing the shortest paths between the winning candidate and the other ones.

Definition 4. Given a complete tournament $G = (\mathcal{C}, \mu)$, a tournament solution S , and a winning candidate $w \in S(G)$, a shortest-path SMS \mathcal{X} for $w \in S(G)$, is an SMS such that for all $c \in \mathcal{C} \setminus \{w\}$, the directed path from w to c in \mathcal{X} is the shortest in G .

Example 5. Looking back at Figure 2, the path from a to d is 1-long in \mathcal{X} and 2-long in \mathcal{Y} as it goes through b . Hence, we will prefer \mathcal{X} with its shorter paths.

For the myopic-score solutions, we argue that it is better to put more emphasis on the good performances of the winner rather than its opponents bad ones. Typically, if a candidate is a Condorcet winner, we would like to show it like \mathcal{I} in Figure 3a. In this perspective, we will prioritize SMSs containing the highest number of won pairwise comparisons by the considered tournament winner.

Definition 5. Given n voters, a partial n -weighted tournament $G = (\mathcal{C}, \mu)$, a tournament solution S , and a winning candidate $w \in S(G)$. Given an SMS $\mathcal{X} = (\mathcal{C}, \mu_{\mathcal{X}})$ for $w \in S(G)$, we define the win count of \mathcal{X} as $WC(\mathcal{X}) = \sum_{c \in \mathcal{C}} \mu_{\mathcal{X}}(w, c)$. \mathcal{X} is a maxwin-SMS if and only if for all SMSs \mathcal{Y} for $w \in S(G)$, we have $WC(\mathcal{X}) \geq WC(\mathcal{Y})$.

4 Computing SMS

In this section we present the technical results of this paper, characterizing the MSs and SMSs for common tournament solutions, and providing polynomial algorithms for their computation when this is possible. All missing proofs can be found in the full version of this paper (Contet, Grandi, and Mengin 2025).

4.1 Top Cycle and Uncovered Set

Top cycle (TC) selects the smallest dominating set of candidates, or, in graph-theoretic perspective, the top strongly-connected component. Here we use an alternative characterization based on paths to prove that all MSs for TC are directed trees.

Proposition 1. *Given a complete tournament $G = (C, E)$, and a winning candidate $w \in \text{TC}(G)$, for all MSs \mathcal{X} for $w \in \text{TC}(G)$, \mathcal{X} is a w -rooted out-tree, i.e., for all $c \in C \setminus \{w\}$, there exists a unique directed path from w to c in \mathcal{X} .*

Proof. Let $\mathcal{X} = (C, E_{\mathcal{X}})$ be an MS for $w \in \text{TC}(G)$. Suppose there exists $c \in C \setminus \{w\}$ such that there is no directed path from w to c in \mathcal{X} . Let \mathcal{Y} be a completion of \mathcal{X} such that w and all the candidates reachable from w in \mathcal{X} lose their missing pairwise comparisons, i.e. such that $\forall (u, v) \in R_w \times C$, $(u, v) \in \mathcal{Y}$ implies $(u, v) \in \mathcal{X}$ or $v \in R_w$ where R_w is the set of reachable candidates from w . Then no new candidate is reachable from w in \mathcal{Y} . Thus, there is no path in \mathcal{Y} from w to c . Hence, $w \notin \text{TC}(\mathcal{Y})$ and $w \notin \text{NW}_{\text{TC}}(\mathcal{X})$ and \mathcal{X} is not an SMS for $w \in \text{TC}(G)$. Contradiction. Hence, for all $c \in C \setminus \{w\}$, there exists a directed path from w to c in \mathcal{X} .

Let us now show that for all $c \in C \setminus \{w\}$, there exists a unique directed path from w to c in \mathcal{X} . Suppose that there exists $c^* \in C \setminus \{w\}$ such that there are two directed paths p' and p'' from w to c^* in \mathcal{X} such that p' ends with $(c', c^*) \in E_{\mathcal{X}}$, p'' ends with $(c'', c^*) \in E_{\mathcal{X}}$ and $c' \neq c''$. For all $c \in C \setminus \{w\}$, we know there exists a directed path p_c from w to c in \mathcal{X} . Let $\mathcal{Y} = (C, E_{\mathcal{Y}})$ be such that $E_{\mathcal{Y}} = E_{\mathcal{X}} \setminus \{(c', c^*)\}$. If p_c does not contain (c'', c^*) , \mathcal{Y} also contains p_c . Else, if p_c contains (c'', c^*) , p_c can be decomposed as $p'' \circ p'_c$ where \circ is the concatenation and p'_c a path from c^* to c . $p' \circ p'_c$ is a directed path from w to c in \mathcal{Y} . For all $c \in C \setminus \{w\}$, there exists a directed path from w to c in \mathcal{Y} , thus $w \in \text{NW}_{\text{TC}}(\mathcal{Y})$ but $\mathcal{Y} \subsetneq \mathcal{X}$ so \mathcal{X} is not an MS for $w \in \text{TC}(G)$. Contradiction. \square

From the common structure shared by MSs for TC, we can show that all MSs and thus SMSs have the same size.

Theorem 2. *Given a complete tournament $G = (C, E)$ with $|C| = m$, and a winning candidate $w \in \text{TC}(G)$, for all MSs \mathcal{X} for $w \in \text{TC}(G)$, we have $|\mathcal{X}| = m - 1$.*

Finally, computing a shortest-path SMS for TC can be achieved in polynomial time with Dijkstra's algorithm (Dijkstra 1959).

Proposition 3. *Given a complete tournament $G = (C, E)$ with $|C| = m$, and a winning candidate $w \in \text{TC}(G)$, there exists an algorithm which computes a shortest-path SMS for $w \in \text{TC}(G)$ in $\mathcal{O}(m^2)$.*

Like TC, the uncovered set (UC) has a characterization based on the existence of directed paths between the winner and the other candidates, except that these paths have a maximum length of two. Hence, the following results can be derived from the previous propositions.

Corollary 4. *Given a complete tournament $G = (C, E)$, and a winning candidate $w \in \text{UC}(G)$, for all MSs \mathcal{X} for $w \in \text{UC}(G)$, \mathcal{X} is a w -rooted out-tree of depth at most 2, i.e., for all $c \in C \setminus \{w\}$, there exists a unique directed path of length at most 2 from w to c in \mathcal{X} .*

Corollary 5. *Given a complete tournament $G = (C, E)$ with $|C| = m$, and a winning candidate $w \in \text{UC}(G)$, for all MSs \mathcal{X} for $w \in \text{UC}(G)$, we have $|\mathcal{X}| = m - 1$.*

Corollary 6. *Given a complete tournament $G = (C, E)$ with $|C| = m$, and a winning candidate $w \in \text{UC}(G)$, there exists an algorithm which computes a shortest-path SMS for $w \in \text{UC}(G)$ in $\mathcal{O}(m^2)$.*

4.2 Weighted Uncovered Set

To construct SMSs for the weighted uncovered set (wUC), one can generalize the alternative characterization of the uncovered set to the case of weighted tournaments given by Döring and Peters (2023).

Lemma 7 (Lemma 2.2 in Döring and Peters 2023). *Given n voters, a complete n -weighted tournament $G = (C, \mu)$, the weighted uncovered set is the nonempty subset of candidates that can reach every other candidate via a directed path (c, c') such that $\mu(c, c') > \mu(c', c)$ or a directed path (c, c', c'') such that $\mu(c, c') > \mu(c'', c')$ in G .*

From Lemma 7, we obtain a result similar to Corollary 4 on the structure of any minimal support for wUC.

Proposition 8. *Given n voters, a complete n -weighted tournament $G = (C, \mu)$, a winning candidate $w \in \text{wUC}(G)$, for all MSs $\mathcal{X} = (C, \mu_{\mathcal{X}})$ for $w \in \text{wUC}(G)$, \mathcal{X} is a w -rooted out-tree of depth at most 2 such that for all losing candidates $c \in C \setminus \{w\}$, $\mu_{\mathcal{X}}(w, c) \geq \lceil (n+1)/2 \rceil$ or there exists a unique c' such that $\mu_{\mathcal{X}}(w, c') + \mu_{\mathcal{X}}(c', c) \geq n + 1$.*

Even if MSs for wUC follow a structure similar to the unweighted case, we now show that deciding whether there exists a MS smaller than a given size is NP-complete.

Theorem 9. *Given $n \geq 2$ voters, a complete n -weighted tournament $G = (C, \mu)$, a winning candidate $w \in \text{wUC}(G)$, and an integer k , determining if there exists an MS of size at most k is NP-complete.*

Proof sketch. To prove NP membership, note that Proposition 8 gives an easy to verify structure to identify if a partial tournament is an MS for wUC.

To prove NP-hardness, we proceed with a reduction from the NP-complete SET COVER problem (Karp 1972). Given an instance of set cover problem with a universe U of p elements and a set S of q subsets of U , we show that there exists a set cover of size less than $k \in \mathbb{N}$ if and only if there exists an MS for wUC of size less than $k' = p + q + k$ for a designated winner in a specific tournament with candidates associated to each element of U and S . \square

Even if computing SMSs for wUC is a hard problem, we provide tight bounds on their size.

Proposition 10. *Given $n \geq 2$ voters, a complete n -weighted tournament $G = (\mathcal{C}, \mu)$ with $|\mathcal{C}| = m \geq 3$, and a winning candidate $w \in \text{wUC}(G)$, for all SMSs \mathcal{X} for $w \in \text{wUC}(G)$, we have $n + m - 2 \leq |\mathcal{X}| \leq (n + 1)(m - 1)$.*

4.3 Maximin

The maximin rule (MM) selects the candidates with the least worst performance in a head-to-head against another alternative. Thus, an MS for MM must guarantee overall good performances for the winner while ensuring that the other candidates perform poorly at least once.

Lemma 11. *Given n voters, a complete n -weighted tournament $G = (\mathcal{C}, \mu)$, and a winning candidate $w \in \text{MM}(G)$ with a maximin score of σ_w , for all MSs $\mathcal{X} = (\mathcal{C}, \mu_{\mathcal{X}})$ for $w \in \text{MM}(G)$, there exists $t \leq \lceil n/2 \rceil$ such that, for all $c \in \mathcal{C} \setminus \{w\}$, $\mu_{\mathcal{X}}(w, c) \geq t$ and there exists $c' \in \mathcal{C}$ such that $\mu_{\mathcal{X}}(c', c) \geq n - t$.*

From the previous characterization, we can derive a closed-form expression for the size of an SMS for MM.

Theorem 12. *Given n voters, a complete n -weighted tournament $G = (\mathcal{C}, \mu)$ with $|\mathcal{C}| = m$, and a winning candidate $w \in \text{MM}(G)$ with a maximin score of σ_w , for all SMSs \mathcal{X} for $w \in \text{MM}(G)$, we have $|\mathcal{X}| = n(m - 1) - t|\{c : c \in \mathcal{C}, \mu(w, c) \geq n - t\}|$ where $t = \min(\sigma_w, \lfloor n/2 \rfloor)$.*

Proof. According to Lemma 11, for each candidate different from w , $n = t + n - t$ pairwise comparisons are needed. However, one can optimize the size of the MS when w wins $n - t$ pairwise comparisons against an opponent c by covering both constraints at once. Indeed, let $\mathcal{X} = (\mathcal{C}, \mu_{\mathcal{X}})$ be an SMS for $w \in \text{MM}(G)$ and $t \leq \lfloor n/2 \rfloor$. For all $c \in \mathcal{C}$ such that $\mu(w, c) \geq n - t$, if we have $\mu_{\mathcal{X}}(w, c) = n - t$, then for all completions \mathcal{X}' of \mathcal{X} , we have $\mu_{\mathcal{X}'}(w, c) \geq n - t \geq t$. For all the other candidates $c \in \mathcal{C}$ such that $\mu(w, c) < n - t$, we take $\mu_{\mathcal{X}}(w, c) = t$ and $\mu_{\mathcal{X}}(c', c) = n - t$ for some candidate $c' \in \mathcal{C}$ such that $\mu(c', c) \geq n - t$. Then for all completions \mathcal{X}' of \mathcal{X} , we have $\mu_{\mathcal{X}'}(w, c) \geq t$ and $\mu_{\mathcal{X}'}(c', c) \geq n - t$. Let $t = \min(\sigma_w, \lfloor n/2 \rfloor)$, for all losing candidates $c \in \mathcal{C} \setminus \{w\}$, either $\mu(w, c) \geq n - t$ and only $n - t$ pairwise comparisons are needed else $n = n - t + t$ comparisons are needed. With $k = |\{c : c \in \mathcal{C}, \mu(w, c) \geq n - t\}|$, we have $|\mathcal{X}| = (n - t)k + n(m - 1 - k) = n(m - 1) - tk$. \square

We provide tight bounds on the size of an SMS for MM.

Corollary 13. *Given n voters, a complete n -weighted tournament $G = (\mathcal{C}, \mu)$ with $|\mathcal{C}| = m$, and a winning candidate $w \in \text{MM}(G)$ with a maximin score of σ_w , for all SMSs \mathcal{X} for $w \in \text{MM}(G)$, we have $\lceil \frac{n}{2} \rceil (m - 1) \leq |\mathcal{X}| \leq n(m - 1)$.*

Computing a maxwin-SMS for MM can be achieved by analyzing the local neighborhood of each candidate.

Proposition 14. *Given n voters, a complete n -weighted tournament $G = (\mathcal{C}, \mu)$ with $|\mathcal{C}| = m$ and a winning candidate $w \in \text{MM}(G)$, there exists an algorithm which computes a maxwin-SMS for $w \in \text{MM}(G)$ in $\mathcal{O}(m^2 \log n)$.*

4.4 Borda and Copeland Rules

The Borda rule (BO) chooses the set of candidates with a maximal total amount of wins across all pairwise comparisons with the other candidates.

We show that the size of an SMS for BO only depends on the Borda score of the winner and its worst performance against the other candidates.

Theorem 15. *Given n voters, a partial n -weighted tournament $G = (\mathcal{C}, \mu)$ with $|\mathcal{C}| = m$, and a winning candidate $w \in \text{BO}(G)$ with a Borda score of σ_w . For all SMSs \mathcal{X} for $w \in \text{BO}(G)$:*

- if $\sigma_w \leq n(m - 1) - \min_{c \neq w} \mu(w, c)$ then $|\mathcal{X}| = (m - 1)(n(m - 1) - \sigma_w)$
- else $|\mathcal{X}| = n(m - 1) - \min(\lfloor n(1 - 1/m) \rfloor, \min_{c \neq w} \mu(w, c))$.

From the characterization of the size of SMSs for BO, we derive tight upper and lower bounds on SMSs size.

Corollary 16. *Given n voters, a partial n -weighted tournament $G = (\mathcal{C}, \mu)$ with $|\mathcal{C}| = m$, and a winning candidate $w \in \text{BO}(G)$. For all SMSs \mathcal{X} for $w \in \text{BO}(G)$, we have $\lceil \frac{n(m-1)^2}{m} \rceil \leq |\mathcal{X}| \leq (m - 1) \lfloor \frac{n(m-1)}{2} \rfloor$.*

We now show that we can compute a maxwin-SMS in polynomial time.

Proposition 17. *Given n voters, a complete n -weighted tournament $G = (\mathcal{C}, E)$ with $|\mathcal{C}| = m$, and a winning candidate $w \in \text{BO}(G)$, there exists an algorithm which computes a maxwin-SMS for $w \in \text{BO}(G)$ in $\mathcal{O}(m^2 \log n)$.*

Finally, we show that all SMSs share a similar structure centered around the winner's wins by bounding the difference in terms of win count between SMSs.

Proposition 18. *Given n voters, an n -weighted tournament $G = (\mathcal{C}, \mu)$, and a winning candidate $w \in \text{BO}(G)$, for all SMSs \mathcal{X}, \mathcal{Y} for $w \in \text{BO}(G)$, we have $|\text{WC}(\mathcal{X}) - \text{WC}(\mathcal{Y})| \leq \max(1, \frac{n}{2})$.*

Since the Copeland rule (CO) coincides with the Borda rule on 1-weighted tournaments, the following results are a special case of the previous results obtained on Borda.

Corollary 19. *Given a complete tournament $G = (\mathcal{C}, E)$ with $|\mathcal{C}| = m$ and a winning candidate $w \in \text{CO}(G)$, for all SMSs \mathcal{X} for $w \in \text{CO}(G)$:*

- if w is a Condorcet winner then $|\mathcal{X}| = m - 1$
- else $|\mathcal{X}| = (m - 1)(m - 1 - \sigma_w)$.

Corollary 20. *Given a complete tournament $G = (\mathcal{C}, E)$ with $|\mathcal{C}| = m$ and a winning candidate $w \in \text{CO}(G)$, for all SMSs \mathcal{X} for $w \in \text{CO}(G)$ $m - 1 \leq |\mathcal{X}| \leq (m - 1) \lfloor \frac{m-1}{2} \rfloor$.*

Corollary 21. *Given a complete tournament $G = (\mathcal{C}, E)$ with $|\mathcal{C}| = m$ and a winning candidate $w \in \text{CO}(G)$, there exists an algorithm which computes a maxwin-SMS for $w \in \text{CO}(G)$ in $\mathcal{O}(m^2)$.*

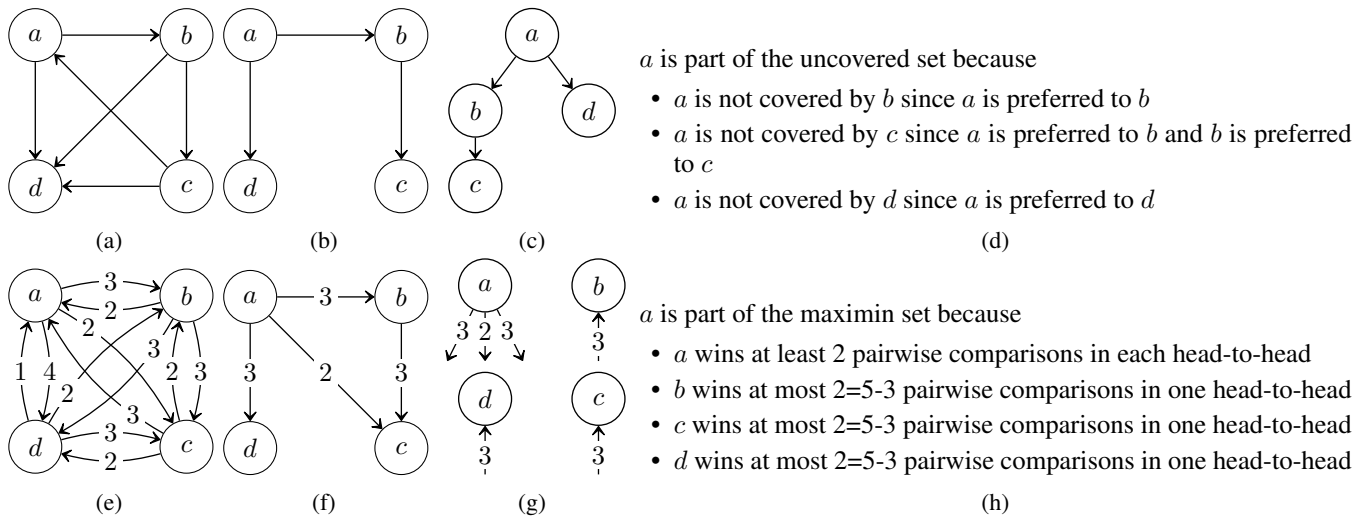


Figure 4: From a tournament (a) and a weighted tournament (e), we compute a smallest minimal support for $a \in UC(G)$ (b) and for $a \in MM(G)$ (f). We identify and visualize their underlying structure as a a -rooted tree (c) or by mean of the out-going edges in the neighborhood of the winning candidate and the in-going edges in those of the losing candidates (g). Finally, we produce textual explanations based on these structures (d) and (h). Examples for the remaining tournament solutions are available in the full version of this paper (Contet, Grandi, and Mengin 2025).

5 Certified User-Friendly Explanations

We now inspect MSs in light of the four Gricean principles underpinning effective explanations that we sketched in the introduction: *quantity*, *quality*, *relation* and *manner* (Grice 1975). The formal grounding of MSs ensure that the reasoning is sound, thus covering *quality*. By definition, MSs provide the minimal information necessary to ensure that a designated candidate is among the election winners. Thus, MSs are informative enough to support our goal and verify both the *quantity* and *relation* principles. Regarding *manner*, or the way information is delivered, it is possible to create explanations containing the smallest amount of information necessary thanks to our algorithms efficiently computing the smallest MSs for five out of the six studied solutions.

However, MSs still suffer from two flaws that limit their explanatory power. First, MSs are (multi-)sets of pairwise comparisons, with a flat structure that makes it hard to gain insights from the unfolding of the tournament solution. Second, their formal nature renders them more suited to be handled by machines or by domain experts than end-users. We propose to address these shortcomings by capitalizing on our structural results for MSs to organize the explanation, and by showcasing how to automatically produce intuitive textual explanation from MSs.

Figure 4 presents two examples of this process. For path-based rules such as UC, MSs can be presented as a rooted out-tree encompassing all directed paths from the winner to the other candidates (Figure 4c), with each path corresponding to a bullet point in the textual explanation (see Figure 4d). For myopic-score rules such as MM, MSs can present the neighborhoods of each candidate (see Figure 4g) from which candidates' scores can be bounded to create a compact textual explanation (see Figure 4h).

6 Conclusions

Minimal supports are compact structures that can serve as certified explanations for tournament solutions. In this paper we have characterized the structure and given the size of smallest minimal supports for well-known tournament solutions, and provided polynomial algorithms for their computation for all but one case. By exploiting our characterizations we have also provided rigorous intuitive visual and textual presentations of a minimal support to demonstrate their applicability as human-oriented explanations.

Minimal supports also have far-reaching connections with several classical problems in computational social choice. In contrast with previous results on the margin of victory (Brill, Schmidt-Kraepelin, and Suksompong 2020; Döring and Peters 2023) for the top cycle, the uncovered set, Copeland, and Borda, we prove that the size of a smallest minimal support is the same for every candidate in the winner set, preventing its use as a refinement of those tournament solutions. From a query complexity perspective, while the top cycle and the uncovered set require $\Omega(m^2)$ queries (Maiti and Dey 2024, Theorem 3.5 & 3.8), we show that, with a perfect oracle, exactly $m - 1$ queries are needed. Although finding the smallest minimal explanation is generally hard, we provide algorithms for the top cycle, the uncovered set, Copeland, maximin and Borda with complexity $\mathcal{O}(m^2 \log n)$, which corresponds to the size of a tournament.

Several key directions merit further exploration: empirically evaluating the usefulness of the presented explanations, studying a broader range of rules, and analyzing structural properties of minimal supports. Finally, in an approach similar to that of De Donder, Le Breton, and Truchon (2000), minimal supports for different rules on the same tournament could be compared from a set-theoretical perspective.

Acknowledgments

The authors thank the reviewers of AAAI26 for their constructive comments and suggestions, which helped improve this paper.

This work is funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. This work is supported by ERC grant 101166894 “Advancing Digital Democratic Innovation” (ADDI).



References

- Amgoud, L.; and Ben-Naim, J. 2022. Axiomatic Foundations of Explainability. In *Proceedings of the 31th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Amgoud, L.; Cooper, M.; and Salim, D. 2024. Axiomatic Characterisations of Sample-based Explainers. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*.
- Arrow, K. J.; Sen, A.; and Suzumura, K. 2010. *Handbook of Social Choice and Welfare*. Elsevier.
- Aziz, H.; Brill, M.; Fischer, F.; Harrenstein, P.; Lang, J.; and Seedig, H. G. 2015. Possible and Necessary Winners of Partial Tournaments. *Journal of Artificial Intelligence Research*, 54: 493–534.
- Beynier, A.; Mailly, J.-G.; Maudet, N.; and Wilczynski, A. 2024. Explaining the Lack of Locally Envy-Free Allocations. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*.
- Boixel, A.; Endriss, U.; and de Haan, R. 2022. A Calculus for Computing Structured Justifications for Election Outcomes. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*.
- Brandt, F.; Brill, M.; and Harrenstein, P. 2016. Tournament Solutions. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, 57–84. Cambridge University Press.
- Brill, M.; Schmidt-Kraepelin, U.; and Suksompong, W. 2020. Refining Tournament Solutions via Margin of Victory. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*.
- Cailloux, O.; and Endriss, U. 2016. Arguing about Voting Rules. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- Chen, L.; Sunny, A. I.; Xu, L.; Xu, S.; Gao, Z.; Lu, Y.; Shi, W.; and Shah, N. 2021. Computational Complexity Characterization of Protecting Elections From Bribery. *Theoretical Computer Science*, 891: 189–209.
- Contet, C.; Grandi, U.; and Mengin, J. 2024. Abductive and Contrastive Explanations for Scoring Rules in Voting. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*.
- Contet, C.; Grandi, U.; and Mengin, J. 2025. Explaining Tournament Solutions with Minimal Supports. arXiv:2509.09312.
- Darwiche, A.; and Hirth, A. 2020. On the Reasons Behind Decisions. In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI)*.
- De Donder, P.; Le Breton, M.; and Truchon, M. 2000. Choosing From a Weighted Tournament. *Mathematical Social Sciences*, 40(1): 85–109.
- Delaunay, J.; Galárraga, L.; Largouët, C.; and Van Berkel, N. 2025. Impact of Explanation Technique and Representation on Users’ Comprehension and Confidence in Explainable AI. *Proceedings of the ACM on Human-Computer Interaction*, 9(2): Article–CSCW113.
- Dijkstra, E. W. 1959. A Note on Two Problems in Connection with Graphs. *Numerische Mathematik*, 50: 269–271.
- Döring, M.; and Peters, J. 2023. Margin of Victory for Weighted Tournament Solutions. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- Eiter, T.; and Gottlob, G. 1995. The Complexity of Logic-Based Abduction. *Journal of the ACM (JACM)*, 42(1): 3–42.
- Elkind, E.; Faliszewski, P.; and Slinko, A. 2009. Swap Bribery. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT)*.
- Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2009. How Hard Is Bribery in Elections? *Journal of Artificial Intelligence Research*, 35: 485–532.
- Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2009. Llull and Copeland Voting Computationally Resist Bribery and Constructive Control. *Journal of Artificial Intelligence Research*, 35: 275–341.
- Fischer, F.; Hudry, O.; and Niedermeier, R. 2016. Weighted Tournament Solutions. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, 85–102. Cambridge University Press.
- Grice, H. P. 1975. Logic and Conversation. In *Speech acts*, 41–58. Brill.
- Grossi, D.; Hahn, U.; Mäs, M.; Nitsche, A.; Behrens, J.; Boehmer, N.; Brill, M.; Endriss, U.; Grandi, U.; Haret, A.; Heitzig, J.; Janssens, N.; Jonker, C. M.; Keijzer, M. A.; Kistner, A.; Lackner, M.; Lieben, A.; Mikhaylovskaya, A.; Murrkannaiah, P. K.; Proietti, C.; Revel, M.; Élise Rouméas; Shapiro, E.; Sreedurga, G.; Swierczek, B.; Talmon, N.; Turcini, P.; Terzopoulou, Z.; and Putte, F. V. D. 2024. Enabling the Digital Democratic Revival: A Research Program for Digital Democracy. arXiv:2401.16863.
- Ignatiev, A.; Narodytska, N.; Asher, N.; and Marques-Silva, J. 2020. On Relating ‘Why?’ and ‘Why Not?’ Explanations. arXiv:2012.11067.

- Karp, R. M. 1972. Reducibility Among Combinatorial Problems. *Complexity of computer computations*, 1: 85–103.
- Konczak, K.; and Lang, J. 2005. Voting Procedures With Incomplete Preferences. In *Proceedings of the Multidisciplinary IJCAI Workshop on Advances in Preference Handling*.
- Kulesza, T.; Stumpf, S.; Burnett, M.; Yang, S.; Kwan, I.; and Wong, W.-K. 2013. Too Much, Too Little, or Just Right? Ways Explanations Impact End Users' Mental Models. In *2013 IEEE Symposium on visual languages and human centric computing*, 3–10. IEEE.
- Liberatore, P. 2005. Redundancy in Logic I: CNF Propositional Formulae. *Artificial Intelligence*, 163(2): 203–232.
- Lipton, P. 1990. Contrastive Explanation. *Royal Institute of Philosophy Supplements*, 27: 247–266.
- Maiti, A.; and Dey, P. 2024. Query Complexity of Tournament Solutions. *Theoretical Computer Science*, 991: 114422.
- Marques-Silva, J. 2022. Logic-Based Explainability in Machine Learning. In *Tutorial Lectures of the 18th International Summer School on Reasoning Web*.
- Miller, T. 2019. Explanation in Artificial Intelligence: Insights From the Social Sciences. *Artificial intelligence*, 267: 1–38.
- Nardi, O.; Boixel, A.; and Endriss, U. 2022. A Graph-Based Algorithm for the Automated Justification of Collective Decisions. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- Peters, D.; Procaccia, A. D.; Psomas, A.; and Zhou, Z. 2020. Explainable Voting. *Advances in Neural Information Processing Systems (NeurIPS)*, 33.
- Shih, A.; Choi, A.; and Darwiche, A. 2018. A Symbolic Approach to Explaining Bayesian Network Classifiers. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Sperner, E. 1928. Ein Satz Über Untermengen Einer Endlichen Menge. *Mathematische Zeitschrift*, 27(1): 544–548.
- Suryanarayana, S. A.; Sarne, D.; and Kraus, S. 2022. Justifying Social-Choice Mechanism Outcome for Improving Participant Satisfaction. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.