

Group Decision Making via Weighted Propositional Logic: Complexity and Islands of Tractability

Gianluigi Greco

Università della Calabria, Italy
ggreco@mat.unical.it

Jérôme Lang

Université Paris-Dauphine, France
lang@lamsade.dauphine.fr

Abstract

We study a general class of multiagent optimization problems, together with a compact representation language of utilities based on weighted propositional formulas. We seek solutions maximizing utilitarian social welfare as well as fair solutions maximizing the utility of the least happy agent. We show that many problems can be expressed in this setting, such as fair division of indivisible goods, some multiwinner elections, or multifacility location. We focus on the complexity of finding optimal solutions, and we identify the tractability boarder between polynomial and NP-hard settings, along several parameters: the syntax of formulas, the allowed weights, as well as the number of agents, propositional symbols, and formulas per agent.

1 Introduction

In a number of application domains, including (some forms of) fair division, voting, or the choice of a common set of items, a decision has to be made by a group of individuals by aggregating their preferences over the available alternatives. In fact, a drastic division exists between *ordinal* settings, where agents express preference relations over alternatives, and *cardinal* settings, where they express utility functions mapping the alternatives to some suitable (typically numerical) scale. In the paper, we focus on the latter settings, and our pragmatic goal is therefore to aggregate such cardinal preferences into a collective utility function.

Now, in group decision-making, the set of possible solutions has often a combinatorial structure: possible allocations of items to agents, coalition structures among agents, binary vectors in multiple referenda, subsets of k candidates in committee elections, etc. The exponential size of the set of solutions implies a tension between expressivity (allowing the agents to express any possible utility function) and elicitation and computation complexity (avoiding the agents to spend hours specifying their preferences, and the computer to spend hours computing the optimal solution).

One common way that sacrifices expressivity but makes elicitation (and often computation) easy consists in assuming that utility functions are *additive*, that is, described only by their values on singletons, the utility of a tuple of values being

then the sum of utilities of the individual values. For instance, when expressing utilities over sets of goods, the utility value given by an agent to a set of goods is the sum of all values she gives to the individual goods in the set. However, assuming additivity implies a huge loss of expressivity, because it does not allow the agents to express *preferential dependencies*. On the other hand, allowing agents to express arbitrary utility functions over a combinatorial set of solutions by listing all solutions together with their utility is clearly unpractical, because it would amount to ask each agent to provide an exponentially large list of values.

A way of reconciling expressivity and complexity is to use a *compact* representation language for representing utility functions. *Weighted propositional logic* is a language of this kind that attracted much attention in the literature: Each individual expresses her preferences as a set of propositional formulas associated with numerical values. Given an interpretation σ assigning a truth value to each variable, the utility of the individual is defined as the *sum*¹ of the values associated with the formulas satisfied by σ . Moreover, in order to preserve the semantics of the application, interpretations might be restricted to those satisfying some given constraints.

Example 1.1. Consider an allocation problem with agents a_1 and a_2 , and three indivisible goods g_1, g_2, g_3 . The setting can be modeled via the set of Boolean variables $\mathcal{V} = \{X_{i,j} \mid i \in \{1, 2\}, j \in \{1, 2, 3\}\}$. An interpretation σ over \mathcal{V} is naturally associated with an allocation, where $X_{i,j}$ true in σ means that a_i receives good g_j . We focus on those interpretations satisfying the formula $\bigwedge_j \bigwedge_{i \neq i'} \neg (X_{i,j} \wedge X_{i',j})$, which constrains each good to be allocated at most to one individual.

Assume a_1 chooses to express her utility function by the set of weighted formulas $\{\langle X_{1,1} \vee (X_{1,2} \wedge X_{1,3}), 3 \rangle, \langle X_{1,1} \wedge X_{1,2}, 2 \rangle\}$, while a_2 has additive preferences, expressed by the set $\{\langle X_{2,1}, 2 \rangle, \langle X_{2,2}, 1 \rangle\}, \langle X_{2,3}, 2 \rangle\}$. Let π be the allocation giving $\{g_1, g_2\}$ to a_1 and $\{g_3\}$ to a_2 . The interpretation corresponding to π is σ_π where the variables evaluating to true are those in $\{X_{1,1}, X_{1,2}, X_{2,3}\}$. The utility of a_1 (resp., a_2) in π is given by $u_1(\sigma_\pi) = 3 + 2 = 5$ (resp., $u_2(\sigma_\pi) = 2$.) \triangleleft

Expressing utilities by weighted formulas is more succinct

¹See [Uckelman and Endriss, 2010] for an alternative approach where the utility is defined as the maximum over the values of the satisfied formulas, and [Lafage and Lang, 2000] for a discussion on (further) possible aggregation functions.

than expressing them directly, and is fully expressive, in the sense that every utility function can be expressed by some set of weighted formulas. Detailed results on the expressivity, succinctness and complexity of various fragments of this language are in [Uckelman *et al.*, 2009]. The succinctness of weighted formulas with respect to other logical representation languages is also discussed in [Coste-Marquis *et al.*, 2004]. Weighted formulas have also been used to express values of coalitions in cooperative games and hedonic games, in so-called *marginal contribution nets* [Jeong and Shoham, 2005; Elkind *et al.*, 2009; Elkind and Wooldridge, 2009], as well as in fair division [Bouveret and Lang, 2008]. Moreover, related languages have been designed for bidding in combinatorial auctions [Boutilier and Hoos, 2001; Nisan, 2006].

Note that, although preference aggregation over combinatorial structures has received a lot of attention these last years, many approaches focus on voting, where preferences are ordinal (see [Lang and Xia, 2015] for a survey); and among general approaches that deal with cardinal preferences, most of them deal with specific application domains (the closest to our work being [E Scoffier *et al.*, 2013]). There are very few existing general (domain-independent) languages allowing for expressing and aggregating cardinal preference. In particular, *GAI networks*, highly related to weighted formulas, have been used by [Gonzales *et al.*, 2008] for expressing group decision making and searching for Pareto-optimal solutions. Logic-based preference representation of collective decision problems has been dealt with in a few papers only [Lafage and Lang, 2000; Uckelman and Endriss, 2010] (see [Endriss, 2011] for a discussion); the main difference with our work is that we address fair optimization in a generic way, and investigate its complexity in depth.

In the paper, we adopt the language of weighted formulas to express individual preferences, and we focus on how these preferences can be aggregated: In the classical *utilitarian* setting, the collective utility (or *utilitarian social welfare*) is the sum of the utilities of the individuals. In *fair* group decision-making, instead, it is more appropriate to consider *egalitarian* (or Rawlsian) social welfare, which is the utility of the least satisfied agent. For instance, under egalitarianism, finding an optimal allocation of indivisible goods to agents is the so-called *Santa Claus problem* [Bezakova and Dani, 2005; Bansal and Sviridenko, 2006].

Example 1.1, continued. The utilitarian social welfare of π is $5 + 2 = 7$. In fact, π is the allocation maximizing the utilitarian social welfare. However, the egalitarian social welfare of π is 2, while the allocation π' giving g_1 to a_1 and $\{g_2, g_3\}$ to a_2 is fairer: its egalitarian social welfare is 3. \triangleleft

Weighted propositional logic can be used in a wide range of collective optimization problems. However, this setting has not been put systematically under the computational lens. In particular, the complexity of the most relevant reasoning problems arising therein and the identification of *islands of tractability* (cf. [Bordeaux *et al.*, 2014; Greco and Scarcello, 2013]), i.e., of classes of instances for which solutions can be efficiently computed, have been unexplored.

The goal of the paper is to fill this gap, by depicting a clear picture of the complexity of collective decision making with

weighted formulas under egalitarianism and by contrasting it with results that hold under utilitarianism. In Section 2, we show that weighted formulas can serve as a general compact representation language for various multiagent optimization problems. In Section 3, we analyze the complexity of deciding the existence of an interpretation guaranteeing some given utilitarian or egalitarian social welfare, and of deciding whether some given interpretation is optimal. In Section 4, starting from the observation that optimal interpretations can be unlikely computed in polynomial time, we define the frontier of tractability by considering all possible combinations for the following restrictions: syntactical restrictions on the propositional language, allowed weights, number of individuals, number of weighted formulas per individual, and maximum number of variables in each formula. In Section 5 and Section 6, the frontier is then charted for the utilitarian and egalitarian social welfare, respectively. Avenues for further research are discussed in Section 7.

2 Formal Framework

Goalbases

Throughout the paper, we assume that a universe \mathcal{V} of variables is given, and that the propositional language \mathcal{P} consists of all formulas built over \mathcal{V} by using the Boolean connectives \wedge , \vee , and \neg , plus the constants \top (true) and \perp (false). For any propositional formula $\varphi \in \mathcal{P}$, $\text{dom}(\varphi)$ denotes the *domain* of φ , i.e., the set of all the variables occurring in it. An *interpretation* $\sigma : \mathcal{W} \rightarrow \{\top, \perp\}$ over $\mathcal{W} \subseteq \mathcal{V}$ is a function assigning a Boolean value to each variable in \mathcal{W} . If $\varphi \in \mathcal{P}$ with $\mathcal{W} \supseteq \text{dom}(\varphi)$, then $\sigma \models \varphi$ means that the interpretation σ is a *model* of φ . A formula φ is *satisfiable* if it has a model.

A *weighted formula* is a pair $\langle \varphi, w \rangle$, where $\varphi \in \mathcal{P}$ is a propositional formula and where $w \in \mathbb{Q}$ is a rational number. A *goalbase* G is a finite set of weighted formulas, whose domain is $\text{dom}(G) = \bigcup_{\langle \varphi, w \rangle \in G} \text{dom}(\varphi)$. For any interpretation σ over a superset of $\text{dom}(G)$, the number $G(\sigma) = \sum_{\langle \varphi, w \rangle \in G \text{ such that } \sigma \models \varphi} w$ is the *value* of σ w.r.t. G .

For any set \mathcal{W} of variables, we denote by $\mathcal{I}(\mathcal{W})$ the set of all interpretations that can be defined over \mathcal{W} .

A *utility function* over \mathcal{W} is a mapping $u : \mathcal{I}(\mathcal{W}) \rightarrow \mathbb{Q}$. Given the function u , we can always build a goalbase G_u with $\text{dom}(G_u) = \mathcal{W}$ and such that $G_u(\sigma) = u(\sigma)$, for each $\sigma \in \mathcal{I}(\mathcal{W})$ [Coste-Marquis *et al.*, 2004; Uckelman *et al.*, 2009].

Utilitarian and Egalitarian Social Welfare

Let \mathcal{G} be a set of goalbases, and let $\text{dom}(\mathcal{G}) = \bigcup_{G \in \mathcal{G}} \text{dom}(G)$ denote the *domain* of \mathcal{G} . Each goalbase $G \in \mathcal{G}$ represents a utility function over $\mathcal{I}(\mathcal{W})$, and we look to define suitable ways to aggregate all these functions into a collective utility function over $\mathcal{I}(\mathcal{W})$. Moreover, the aggregation process is often subject to constraints emerging from the application, which can be naturally modeled (again) as formulas in \mathcal{P} that have to be satisfied by the candidate interpretations.

Definition 2.1. A (*group decision-making*) *scenario* is a pair (\mathcal{G}, Γ) where \mathcal{G} is a set of goalbases and Γ (the constraint) is a satisfiable propositional formula in \mathcal{P} . An interpretation $\sigma \in \mathcal{I}(\text{dom}(\mathcal{G}))$ is *feasible* (in (\mathcal{G}, Γ)) if $\sigma \models \Gamma$. \square

For instance, in Example 1.1, the fact that each good can be allocated at most to one individual is encoded via the constraint $\Gamma = \bigwedge_j \bigwedge_{i \neq i'} \neg (X_{i,j} \wedge X_{i',j})$. Therefore, the pair $(\{G_1, G_2\}, \Gamma)$ where G_1 and G_2 are the goalbases of agents a_1 and a_2 , respectively, is a group decision-making scenario formalizing the allocation problem introduced there.

A scenario (\mathcal{G}, Γ) is *positive* (resp., *negative*) if every weighted formula occurring in the goalbases in \mathcal{G} is associated with a positive (resp., negative) weight—w.l.o.g., weighted formulas of the form $\langle \varphi, 0 \rangle$ can be removed.

Let σ be any feasible interpretation in $\mathcal{I}(\text{dom}(\mathcal{G}))$. Then, the *utilitarian social welfare* of σ is the value $\text{sw}(\sigma) = \sum_{G \in \mathcal{G}} G(\sigma)$. We say that σ is *sw-optimal* if it has the maximum utilitarian social welfare over all feasible interpretations in $\mathcal{I}(\text{dom}(\mathcal{G}))$. The set of all sw-optimal interpretations is denoted by $\text{SW}^*\text{-INT}(\mathcal{G}, \Gamma)$, and their utilitarian social welfare is denoted by $\text{SW}^*\text{-VAL}(\mathcal{G}, \Gamma)$.

Now, the intuition in the exposition that follows is to look for interpretations that are not “too far” from the optimum values that can be achieved when optimizing each of the goalbases independently on the others. The definition below uses a normalization mechanism to uniformly deal with utilities defined over different scales. We will see that, whenever preferences are already normalized, as it is often assumed in fair optimization (see, e.g., [Escoffier *et al.*, 2013]), the definition reduces to a classical *max-min* fairness criterion.

Definition 2.2. Let (\mathcal{G}, Γ) be a group decision-making scenario, let $\alpha \in \mathbb{Q}$, and let σ be a feasible interpretation. We say that σ is α -*fair* (in (\mathcal{G}, Γ)) if, for each $G \in \mathcal{G}$,

$$\text{SW}^*\text{-VAL}(\{G\}, \Gamma) - G(\sigma) \leq (1 - \alpha)(\text{SW}^*\text{-VAL}(\{G\}, \Gamma) - \text{MIN}(G, \Gamma))$$

where $\text{MIN}(G, \Gamma)$ denotes the minimum value $G(\sigma')$ achieves over all possible feasible interpretations σ' . \square

Note that every feasible interpretation σ is 0-fair. Then, the *egalitarian value* of σ , denoted by $\text{EG}(\sigma)$, is the maximum value α for which σ is α -fair. Equivalently,

$$\text{EG}(\sigma) = \min_i \frac{G_i(\sigma) - \text{MIN}(G_i, \Gamma)}{\text{SW}^*\text{-VAL}(\{G_i\}, \Gamma) - \text{MIN}(G_i, \Gamma)}.$$

Note that $\text{EG}(\sigma) \leq 1$. In particular, if $\text{EG}(\sigma) = 1$, then σ is an interpretation over which every goalbase can achieve its optimal value. We say that σ is *EG-optimal* if it has the maximum egalitarian value over all possible feasible interpretations. The set of all EG-optimal interpretations is denoted by $\text{EG}^*\text{-INT}(\mathcal{G}, \Gamma)$, and $\text{EG}^*\text{-VAL}(\mathcal{G}, \Gamma)$ is their egalitarian value.

Example 1.1, continued. Each agent can get all objects if she were alone; hence, $\text{SW}^*\text{-VAL}(\{G_1\}) = \text{SW}^*\text{-VAL}(\{G_2\}) = 5$. Moreover, $\text{MIN}(G_1, \Gamma) = \text{MIN}(G_2, \Gamma) = 0$. Thus, an interpretation σ is α -fair if, and only if, $G_1(\sigma) \geq \alpha \times 5$ and $G_2(\sigma) \geq \alpha \times 5$. Accordingly, the interpretation $\sigma_{\pi'}$, where the variables evaluating to true are those in $\{X_{1,1}, X_{2,2}, X_{2,3}\}$, is such that $\text{EG}(\sigma_{\pi'}) = \min\{\frac{3}{5}, \frac{3}{5}\} = \frac{3}{5}$. Note that $\sigma_{\pi'}$ is EG-optimal, and we have $\text{EG}^*\text{-VAL}(\{G_1, G_2\}, \Gamma) = \frac{3}{5}$. \triangleleft

Examples (in the Egalitarian Setting)

Weighted formulas are easily seen to apply in the utilitarian setting (cf. [Uckelman *et al.*, 2009]). We next show that they apply to several problems in the egalitarian setting, too.

For the sake of exposition, we shall use the propositional language extended by *cardinality formulas*. In particular, the formula $\geq k : \varphi_1, \dots, \varphi_p$ is satisfied by an interpretation σ if $|\{\varphi_i \mid i \in \{1, \dots, p\}, \sigma \models \varphi_i\}| \geq k$ holds—similarly for $\leq k : \varphi_1, \dots, \varphi_p$ and for $= k : \varphi_1, \dots, \varphi_p$.²

Fair division of indivisible goods. As we have seen in Example 1.1, fair division of indivisible goods can be naturally expressed in our general framework. Variations of the problem involving, for instance, complex constraints over allocations (such as when some goods can be assigned to more than one agent and some not, etc.) can be expressed as well.

Minimax approval voting (e.g., [Brams *et al.*, 2007]): We have to build a committee of exactly k persons among m candidates from the set $\mathcal{K} = \{P_1, \dots, P_m\}$. Each of n voters approves a subset $A_i \subseteq \mathcal{K}$. A set $C \subseteq \mathcal{K}$, with $|C| = k$, is a minimax approval committee if it minimizes the value $\max_i |C \setminus A_i \cup A_i \setminus C|$. The problem can be modeled as a scenario $(\{G_1, \dots, G_n\}, \Gamma)$, where candidates in $\{P_1, \dots, P_m\}$ are viewed as variables. Intuitively, for any interpretation σ , $\sigma(P_j) = \top$ (resp., \perp) means that P_j is included (resp., not included) in the committee. The scenario is such that for each voter i , $G_i = \{\langle P_j, 1 \rangle \mid P_j \in A_i\} \cup \{\langle \neg P_j, 1 \rangle \mid P_j \notin A_i\}$, and Γ is $=k : P_1, \dots, P_m$. So, interpretations σ maximizing $\text{EG}(\sigma)$ one-to-one correspond to minimax approval committees.

Multifacility location (e.g., [Elzinga *et al.*, 1976]): Let $\mathcal{L} = \{L_1, \dots, L_m\}$ be a set of locations, and let $u_i(L_j)$ be the reward for agent i when using a facility located in L_j . We want to open at most k facilities, and each agent will use the one providing the best reward to her. The problem can be modeled as a scenario $(\{G_1, \dots, G_n\}, \Gamma)$ defined over the variables in \mathcal{L} . For any interpretation σ , $\sigma(L_j) = \top$ (resp., $\sigma(L_j) = \perp$) means that a facility is opened (resp., not opened) at location L_j . For each agent a_i , let us first order the locations according to their nonincreasing order of u_i : $u_i(L_{i_1}) \geq \dots \geq u_i(L_{i_m})$, with $\{i_1, \dots, i_m\} = \{1, \dots, m\}$. Then, we define:

$$G_i = \left\{ \begin{array}{l} \langle L_{i_1}, u_i(i_1) \rangle, \langle \neg L_{i_1} \wedge L_{i_2}, u_i(i_2) \rangle, \dots, \\ \langle \neg L_{i_1} \wedge \dots \wedge \neg L_{i_{m-1}} \wedge L_{i_m}, u_i(i_m) \rangle \end{array} \right\}.$$

Hence, $G_i(\sigma)$ is the best reward of agent a_i derived from using any of the facilities that are opened (according to σ). Finally, Γ is the constraint $\leq k : L_1, \dots, L_m$. Note that EG-optimal interpretations one-to-one correspond to solutions where facilities are opened as to maximize the (normalized) reward of the least satisfied agent. Similar problems where each agent enjoys the best item in the selected subset are *budgeted social choice* [Lu and Boutilier, 2011] and *full proportional representation* [Chamberlin and Courant, 1983; Monroe, 1995; Procaccia *et al.*, 2008; Betzler *et al.*, 2013; Skowron *et al.*, 2013].

Fair group knapsack (e.g., [Shachnai and Tamir, 2001]): The setting is similar to fair division of indivisible goods, except the constraint no longer prescribing that goods are non-sharable, but rather enforcing a bound on the number (or volume) of goods that can be put in the knapsack.

²Cardinality formulas can be rewritten in standard propositional logic, by introducing additional variables but without an exponential blow-up in their size [van Hentenryck and Deville, 1991].

3 Maximizing Social Welfare

Given $X \in \{\text{SW}, \text{EG}\}$ and a scenario (\mathcal{G}, Γ) provided as input, we focus on the following two decision problems:

X-VAL-CHECK: Given $\gamma \in \mathbb{Q}$, does $X^*\text{-VAL}(\mathcal{G}, \Gamma) \geq \gamma$ hold?

X-INT-CHECK: Given $\sigma \in \mathcal{I}(\mathcal{V})$, does $\sigma \in X^*\text{-INT}(\mathcal{G}, \Gamma)$ hold?

We start with the utilitarian social welfare, where intractability is established over *single-agent* scenarios (\mathcal{G}, Γ) , i.e., with $|\mathcal{G}| = 1$, and *without constraints* (i.e., $\Gamma = \top$).

Theorem 3.1. SW-VAL-CHECK (resp., SW-INT-CHECK) is **NP-complete** (resp., **co-NP-complete**). Hardness holds on positive and negative single-agent scenarios without constraints.

Moving to the egalitarian social welfare, we have to observe that according to Definition 2.2, even just checking whether some given interpretation is α -fair requires the computation of the maximum *utilitarian* social welfare associated with each of the individual goalbases. We show below that this problem is $\mathbf{F}\Delta_2^{\mathbf{P}}$ -hard. For completeness, recall that $\mathbf{F}\Delta_2^{\mathbf{P}}$ (resp., $\Delta_2^{\mathbf{P}}$) is the set of all computation (decision) problems solvable in polynomial time by using a **NP** oracle whose cost is assumed to be unitary—see, e.g., [Papadimitriou, 2003].

Proposition 3.2. Computing the maximum utilitarian social welfare is $\mathbf{F}\Delta_2^{\mathbf{P}}$ -hard, even in single agent scenarios.

Proof Sketch. Consider the $\Delta_2^{\mathbf{P}}$ -complete problem MAX-ASG-ODD [Wagner, 1987]: given a satisfiable formula Σ over $\{X_1, \dots, X_n\}$, does X_n evaluates to true in the lexicographically greatest model (w.r.t. the order $\{X_1 > \dots > X_n\}$)? The result follows by building in polynomial time $(\{G\}, \Sigma)$, where $G = \{\langle X_1, 2^{n-1} \rangle, \dots, \langle X_n, 2^0 \rangle\}$. \square

The result implies that normalization in Definition 2.2 typically determines a complexity increase. Take for instance the problem EG-VAL-CHECK. *Without normalization*, the problem is **NP-complete** (membership in **NP** is clear, since it suffices to guess an interpretation σ and check that $G_i(\sigma) \geq \gamma$, for each i). However, *with normalization*, it climbs up to $\Delta_2^{\mathbf{P}}$.

Theorem 3.3. Given a rational number γ , checking whether $\text{EG}^*\text{-VAL}(\mathcal{G}, \Gamma) \geq \gamma$ holds is $\Delta_2^{\mathbf{P}}$ -complete.

To circumvent the issue, we shall hereinafter reason on scenarios where values are explicitly normalized. Formally, a scenario (\mathcal{G}, Γ) is *normalized* if $\text{SW}^*\text{-VAL}(\{G\}, \Gamma) = 1$ and $\text{MIN}(G, \Gamma) = 0$, for each $G \in \mathcal{G}$. Over normalized scenarios (\mathcal{G}, Γ) , fair optimization reduces to a max-min approach:

$$\text{EG}^*\text{-VAL}(\mathcal{G}, \Gamma) = \max_{\sigma \in \mathcal{I}(\text{dom}(\mathcal{G})), \sigma \models \Gamma} \left(\min_{G \in \mathcal{G}} G(\sigma) \right).$$

Note that any scenario (\mathcal{G}, Γ) can be transformed into an equivalent normalized one $(\hat{\mathcal{G}}, \Gamma)$ by removing from \mathcal{G} goalbases G' such that $\text{SW}^*\text{-VAL}(\{G'\}, \Gamma) - \text{MIN}(G', \Gamma) = 0$, and by replacing each remaining goalbase G with the modified goalbase containing, for each $\langle \varphi, c \rangle \in G$, the two weighted formulas $\langle \varphi, c/H \rangle$ and $\langle \top, -\text{MIN}(G, \Gamma)/H \rangle$, where $H = \text{SW}^*\text{-VAL}(\{G\}, \Gamma) - \text{MIN}(G, \Gamma)$. In fact, to constructively apply the normalization procedure, a bottleneck to be faced is the computation of the utilitarian social welfare. That is, the source of complexity underlying the $\Delta_2^{\mathbf{P}}$ -hardness in Theorem 3.3 is now made “explicit” in this pre-processing

	arbitrary		positive	negative
	SW	EG	SW	SW
VAL-CHECK	NP-c	NP-c	NP-c	NP-c
INT-CHECK	co- NP-c	co- NP-c	co- NP-c	co- NP-c

Figure 1: Summary of complexity results. Under the egalitarian semantics, scenarios are normalized (and non-trivial).

phase. Note that if $\hat{\mathcal{G}} = \emptyset$, then any feasible interpretation for (\mathcal{G}, Γ) is trivially EG-optimal. Otherwise, the equivalence of (\mathcal{G}, Γ) and $(\hat{\mathcal{G}}, \Gamma)$ is stated below—the proof is simple.

Proposition 3.4. If $\hat{\mathcal{G}} \neq \emptyset$, then $\text{EG}^*\text{-VAL}(\hat{\mathcal{G}}, \Gamma) = \text{EG}^*\text{-VAL}(\mathcal{G}, \Gamma)$, $\text{EG}^*\text{-INT}(\hat{\mathcal{G}}, \Gamma) = \text{EG}^*\text{-INT}(\mathcal{G}, \Gamma)$, and $(\hat{\mathcal{G}}, \Gamma)$ is normalized.

In the following, we assume that any scenario (\mathcal{G}, Γ) provided as input to problems under egalitarianism is normalized, and that \mathcal{G} is non-trivial (i.e., $|\mathcal{G}| > 1$, for otherwise egalitarianism is immaterial). The counterpart of Theorem 3.1 is stated below.

Theorem 3.5. EG-VAL-CHECK (resp., EG-INT-CHECK) is **NP-complete** (resp., **co-NP-complete**).

We leave the section by noticing that after normalization, complexity results derived for utilitarianism and for egalitarianism coincide—see Figure 1. In fact, under egalitarianism (over normalized scenarios), “negative” scenarios are not possible and arbitrary scenarios coincide with “positive” ones.

4 Charting the Frontier of Tractability

The results derived so far suggest that the problems of interest are unlikely to be solvable in polynomial time, even by considering restrictions on the weights associated with the various formulas. This motivates to identify islands of tractability by considering further restrictions, which we do now.

We start by defining different restrictions on formulas and weights. Consider the language $\mathcal{L}_{\{\wedge, \vee, \neg\}}$ consisting of all propositional formulas φ built according to the following grammar: $\varphi ::= X \mid \neg X \mid (\varphi \wedge \dots \wedge \varphi) \mid (\varphi \vee \dots \vee \varphi)$, where X is any variable in \mathcal{V} . If $C \subseteq \{\wedge, \vee, \neg\}$ is a set of Boolean connectives, then we write \mathcal{L}_C to denote the set of all the formulas in $\mathcal{L}_{\{\wedge, \vee, \neg\}}$ that do not contain symbols in $C \setminus \{\wedge, \vee, \neg\}$. Note that we assume that all formulas are in Negation Normal Form, that is, that negation applies only over variables (if it could apply over general subformulas, then \wedge would be expressible from \vee and vice versa).

Let $h_1, h_2, h_3 \in \{1, c, \infty\}$, let $S \subseteq \{+, -\}$ with $|S| \geq 1$, and let $\mathcal{GB}_{C,S}[h_1, h_2, h_3]$ be the set of scenarios (\mathcal{G}, Γ) where:

- Γ as well as all formulas in \mathcal{G} are taken from \mathcal{L}_C ;
- if $S = \{+\}$ (resp., $S = \{-\}$), then (\mathcal{G}, Γ) is positive (resp., negative); if $S = \{+, -\}$ then no restriction is imposed on the weights associated with the formulas;
- if $h_1 = 1$ (resp., $h_1 = c$), then $|\mathcal{G}| = 1$ (resp., $|\mathcal{G}|$ is bounded by a fixed constant); if $h_1 = \infty$, then no bound is required over $|\mathcal{G}|$;
- if $h_2 = 1$ (resp., $h_2 = c$), then $|G_i| = 1$ (resp., $|G_i|$ is bounded by a fixed constant), for each $G_i \in \mathcal{G}$; if $h_2 = \infty$, then no bound is required over any $|G_i|$; and

		$h_3 \Rightarrow$		
		1	c	∞
h_2 \Downarrow	1	P	P	NP-h^o
	c	P	P	NP-h^o
	∞	P	NP-h*	NP-h[†]

*FrT = { ({ \wedge , \vee }, {+}), ({ \wedge , \vee }, {-}), ({-}, {+, -}) }
^oFrT = { ({ \wedge , \vee }, {+}), ({ \wedge , \neg }, {+, -}), ({ \vee , \neg }, {+, -}) }
[†]FrT = { ({ \wedge , \vee }, {+}), ({ \wedge }, {-}), ({ \vee }, {-}), ({-}, {+, -}) }

		$h_3 \Rightarrow$		
		1	c	∞
h_2 \Downarrow	1	P	P	NP-h^o
	c	P	P	NP-h^o
	∞	NP-h*	NP-h*	NP-h*

$h_1 = c$

		$h_3 \Rightarrow$		
		1	c	∞
h_2 \Downarrow	1	P	NP-h[‡]	NP-h[‡]
	c	NP-h*	NP-h*	NP-h*
	∞	NP-h*	NP-h*	NP-h*

$h_1 = \infty$

*FrT = { ({ \wedge , \vee }, {+}), ({ \wedge , \vee , \neg }, {-}) }
^oFrT = { ({ \wedge , \vee }, {+}), ({ \wedge , \neg }, {+, -}), ({ \vee , \neg }, {+, -}), ({ \wedge , \vee , \neg }, {-}) }
[‡]FrT = { ({ \wedge , \vee }, {+, -}), ({ \wedge , \neg }, {+, -}), ({ \wedge , \vee , \neg }, {-}) }
[†]FrT = { ({ \wedge , \vee }, {+, -}), ({ \wedge , \neg }, {+, -}), ({ \wedge , \vee , \neg }, {-}) }

Figure 2: Summary of results: Individual decision-making (left) and egalitarian social welfare (right). Entries marked as “P” correspond to tractable settings (without any restriction). Instead, for the NP-hard scenarios, the tractability frontier is reported.

- if $h_3 = 1$ (resp., $h_3 = c$), then $|\text{dom}(\varphi_i)| = |\text{dom}(\Gamma)|$ are equal to 1 (resp., are bounded by a fixed constant) for each $G_i \in G$ and for $\langle \varphi_i, w_i \rangle \in G_i$; if $h_3 = \infty$, then no bound is required over $|\text{dom}(\varphi_i)|$ and $|\text{dom}(\Gamma)|$.

The basic computational task we shall consider is the problem of computing an X-optimal interpretation restricted on $\mathcal{GB}_{C,S}[h_1, h_2, h_3]$, which is hereinafter denoted as X-FIND $_{C,S}[h_1, h_2, h_3]$.³ To analyze its complexity, we define a concept of *frontier of tractability*, which allows us to express more succinctly our results, by avoiding to get lost in the very large number of different combinations we shall consider.

Definition 4.1. The *frontier of tractability* of $X \in \{\text{SW}, \text{EG}\}$ w.r.t. $h_1, h_2, h_3 \in \{1, c, \infty\}$, denoted by X-FrT $[h_1, h_2, h_3]$, is the *minimal* set of pairs (C, S) with $C \subseteq \{\vee, \wedge, \neg\}$ and $\emptyset \subset S \subseteq \{+, -\}$ such that:

- $\forall (C, S) \in \text{X-FrT}[h_1, h_2, h_3]$, X-FIND $_{C,S}[h_1, h_2, h_3]$ is in **P**;
- $\forall (C', S')$ with $C' \subseteq \{\wedge, \vee, \neg\}$ and $S' \subseteq \{+, -\}$ such that there is no pair $(C, S) \in \text{X-FrT}[h_1, h_2, h_3]$ with $C \cup S \supseteq C' \cup S'$, X-FIND $_{C',S'}[h_1, h_2, h_3]$ is **NP-hard**. \square

Note that the notion precisely captures the intuition that the given pairs mark the boundary between tractable and intractable settings. The following two sections are then devoted to present the frontier of tractability for the utilitarian and egalitarian social welfare, respectively.

5 Results for Utilitarian Social Welfare

We start the analysis with the case $h_1 = 1$, which is the classical setting of individual decision-making. The use of goalbases has been largely explored in this case, but the frontier of tractability has not been charted according to the various qualitative parameters we are considering.

Theorem 5.1. *Results for individual optimization are summarized in Figure 2.*

Note that the frontier emerged to be rather fragmented. This needed to establish a large number of NP-hardness and tractability results in order to chart it. Since we cannot report all proofs here, we focus below on discussing two scenarios that are representative of the techniques we have adopted. First, we consider a tractable class.

³If X-FIND $_{C,S}[h_1, h_2, h_3]$ is in **P**, then we can answer X-VAL-CHECK and X-INT-CHECK in polynomial time, too.

Lemma 1. SW-FIND $_{\{\wedge, \vee\}, \{+\}}[1, \infty, c]$ is in **P**.

Proof. Let (\mathcal{G}, Γ) be in $\mathcal{GB}_{\{\wedge, \vee\}, \{+\}}[1, \infty, c]$, and observe that the constraint Γ contains a constant number of variables only. So, we can proceed as follows. First, we can explicitly enumerate all possible assignments σ_Γ for these variables. For each assignment σ_Γ such that $\sigma_\Gamma \models \Gamma$, we can solve the scenario $(\mathcal{G}[\sigma_\Gamma], \top)$, where $\mathcal{G}[\sigma_\Gamma]$ is derived from \mathcal{G} by just replacing each variable in the domain of σ_Γ with \top or \perp depending on its truth value in σ_Γ . Eventually, over all SW-optimal interpretations computed over the scenarios of this kind, we return the one with the maximum associated possible SW-value. In fact, all rules in $(\mathcal{G}[\sigma_\Gamma], \top)$ are positive and negation is not allowed. Thus, an optimal interpretation for $(\mathcal{G}[\sigma_\Gamma], \top)$ is one where all variables are mapped to true (see Theorem 5.6 in [Uckelman *et al.*, 2009]). \square

The proof for an intractable scenario is now illustrated.

Lemma 2. SW-FIND $_{\{\vee, \neg\}, \{-\}}[1, \infty, c]$ is in **NP-hard**.

Proof. Let $\mathcal{I} = \{I_1, \dots, I_n\}$ be a collection of items, and $\mathcal{S} = \{S_1, \dots, S_m\}$ a set of subsets of \mathcal{I} . An exact cover (over \mathcal{I} and \mathcal{S}) is a set $S' \subseteq \mathcal{S}$ such that I_i belongs to exactly one element of S' for all $i \in \{1, \dots, n\}$. Deciding whether there is an exact cover is known to be NP-hard, even if $|S_j| = 3$ holds for each $j \in \{1, \dots, m\}$, and if each element I_i , with $i \in \{1, \dots, n\}$, belongs precisely to three different sets in \mathcal{S} . Given \mathcal{I} and \mathcal{S} , we build the goalbase $G = G_1 \cup \dots \cup G_n$, where each set G_i is defined as follows. Elements in \mathcal{S} are transparently viewed as variables. For each item $I_i \in \mathcal{I}$, if S_{i_1}, S_{i_2} , and S_{i_3} are the sets where I_i occurs, then define $G_i = \{\langle \neg S_{i_k} \vee S_{i_h} \vee S_{i_f}, -1 \rangle \mid h, k, f \in \{1, 2, 3\} \text{ with } h \neq k, h \neq f, f \neq k\}$. Let σ be any interpretation over $\{S_{i_1}, S_{i_2}, S_{i_3}\}$. Observe that $G_i(\sigma) = -4$ (resp., $G_i(\sigma) = -6$) if, and only if, exactly one of the variables in the domain evaluates true (resp., either none of the variables or more than one variable evaluate true). Therefore, there is an exact cover if, and only if, $\text{SW}^*\text{-VAL}(\{G\}) = -4 \times n$. \square

Moving from $h_1 = 1$ to $h_1 \in \{c, \infty\}$, we just observe that the optimization of a set \mathcal{G} of goalbases is equivalent to the optimization of $G = \{\langle \varphi, w \rangle \mid \exists G_i \in \mathcal{G} \text{ such that } \langle \varphi, w \rangle \in G_i\}$. Thus, the following can be established.

Corollary 5.2. *For each $h_3 \in \{1, c, \infty\}$, it holds that*

- (1) SW-FrT $[c, 1, h_3] = \text{SW-FrT}[c, c, h_3] = \text{SW-FrT}[1, c, h_3]$;
- (2) SW-FrT $[c, \infty, h_3] = \text{SW-FrT}[1, \infty, h_3]$; and
- (3) SW-FrT $[\infty, h_2, h_3] = \text{SW-FrT}[1, \infty, h_3]$, $h_2 \in \{1, c, \infty\}$.

6 Results for Egalitarian Social Welfare

We now focus on the egalitarian social welfare.

Theorem 6.1. *Results for the egalitarian social welfare are summarized in Figure 2.*

Recall that, under egalitarianism, we assume that scenarios are given which are normalized and non-trivial. Accordingly, the classes $\mathcal{GB}_{\{\wedge, \vee, \neg\}, \{-\}}[h_1, h_2, h_3]$ are considered as (trivially) tractable, for notational uniformity. For the remaining classes in Figure 2, we discuss representative cases. We start with the proof for a tractability frontier with $h_2 = \infty$.

Lemma 3. *For each $h_3 \in \{1, c, \infty\}$, $\text{EG-FrT}[c, \infty, h_3] = \text{EG-FrT}[\infty, \infty, h_3] = \{(\{\wedge, \vee\}, \{+\}), (\{\wedge, \vee, \neg\}, \{-\})\}$.*

Proof Sketch. We observe that whenever $C = \{\wedge, \vee\}$ and (\mathcal{G}, Γ) is positive, then the interpretation mapping all variables to true is in $\text{EG}^*\text{-VAL}(\mathcal{G}, \Gamma)$. It remains to focus on the class $\mathcal{GB}_{C,S}[c, \infty, 1]$, by distinguishing two cases.

(1) $C \supseteq \{\neg\}$ and $S = \{+\}$. We show that the setting is NP-hard with a reduction from the well-known PARTITION problem: given a multiset $S = \{s_1, s_2, \dots, s_n\}$ of natural numbers, is there a partition of S into two multisets S_1 and S_2 such $\sum_{p \in S_1} p = \sum_{q \in S_2} q$? Indeed, given this instance, we can consider two goalbases G_1 and G_2 built over the variables $\{X_1, \dots, X_n\}$ and such that $G_1 = \bigcup_{s_i \in S} \{\langle X_i, \frac{s_i}{M} \rangle\}$ and $G_2 = \bigcup_{s_i \in S} \{\langle \neg X_i, \frac{s_i}{M} \rangle\}$, with $M = \sum_{s_i \in S} s_i$. Note that $\text{SW}^*\text{-VAL}(\{G_1\}, \top) = \text{SW}^*\text{-VAL}(\{G_2\}, \top) = 1$. Moreover, $(\{G_1, G_2\}, \top)$ is normalized, and $\text{EG}^*\text{-VAL}(\{G_1, G_2\}, \top) = \frac{1}{2}$ if, and only if, S is a positive instance of PARTITION.

(2) $C \supseteq \{\}$ and $S = \{+, -\}$. We exhibit another reduction from PARTITION. Indeed, let us build the setting with the two goalbases $G'_1 = \bigcup_{s_i \in S} \{\langle X_i, \frac{s_i}{M} \rangle\}$ and $G'_2 = \bigcup_{s_i \in S} \{\langle X_i, -\frac{s_i}{M} \rangle\} \cup \{\langle \top, 1 \rangle\}$. The scenario $(\{G'_1, G'_2\}, \top)$ is normalized, and it can be again checked that $\text{EG}^*\text{-VAL}(\{G'_1, G'_2\}, \top) = \frac{1}{2}$ if, and only if, S is a positive instance of PARTITION. \square

Note that the island of tractability identified in the above result is rather small. Good news comes instead for $h_2 = c$.

Lemma 4. *The frontier $\text{EG-FrT}[c, c, \infty]$ is $\{(\{\wedge, \vee\}, \{+\}), (\{\wedge, \neg\}, \{+, -\}), (\{\vee, \neg\}, \{+, -\}), (\{\wedge, \vee, \neg\}, \{-\})\}$*

Proof Sketch. (*Tractability Results*) Observe that for the class $\mathcal{GB}_{\{\wedge, \vee\}, \{+\}}[c, c, \infty]$, the interpretation mapping all variables to true is optimal. Therefore, we need only to focus on the pairs (C, S) in $\{(\{\wedge, \neg\}, \{+, -\}), (\{\vee, \neg\}, \{+, -\})\}$.

Let (C, S) be one of these two pairs, and let (\mathcal{G}, Γ) be in $\mathcal{GB}_{C,S}[c, c, \infty]$. Consider an algorithm that iterates over each goalbase G_i in $\{G_1, \dots, G_k\} = \mathcal{G}$ and over each subset G'_i of G_i . At each iteration, we build the formula

$$\Phi_{G'_i} = \bigwedge_{\langle \varphi_i, w_i \rangle \in G'_i} \varphi_i \wedge \bigwedge_{\langle \bar{\varphi}_i, \bar{w}_i \rangle \in G_i \setminus G'_i} \neg(\bar{\varphi}_i).$$

Then, for each set $\{G'_1, \dots, G'_k\}$ with $G'_i \subseteq G_i$, we look for an interpretation $\sigma_{\{G'_1, \dots, G'_k\}}$ such that $\sigma_{\{G'_1, \dots, G'_k\}} \models \Phi_{\{G'_1, \dots, G'_k\}}$, where $\Phi_{\{G'_1, \dots, G'_k\}} = \Gamma \wedge \Phi_{G'_1} \wedge \dots \wedge \Phi_{G'_k}$.

In fact, the weighted formulas included in G_i and satisfied by $\sigma_{\{G'_1, \dots, G'_k\}}$ are precisely those in G'_i (and the associated egalitarian value can be immediately computed given

$\sigma_{\{G'_1, \dots, G'_k\}}$). Eventually, the feasible interpretation with the best possible egalitarian value will be returned. Now, observe that the computation requires a polynomial number of steps, since h_1 and h_2 are bounded by some constant. Moreover, we can decide in polynomial time whether $\Phi_{\{G'_1, \dots, G'_k\}}$ is satisfiable and that, whenever this is the case, then a satisfying assignment can be actually computed in polynomial time, too. Indeed, just note that $\Phi_{\{G'_1, \dots, G'_k\}}$ can be easily rewritten as a Boolean formula in conjunctive normal form, where only a constant number of clauses can contain more than one literal.

(*Hardness Results*) Consider first the case where $C = \{\wedge, \vee, \neg\}$ and $S = \{+\}$. Given a Boolean formula φ , consider a variable X not in $\text{dom}(\varphi)$. Consider the goalbases $G_\varphi = \{\langle X \wedge \varphi, 1 \rangle, \langle \neg X, 1 \rangle\}$ and $\bar{G}_\varphi = \{\langle X, 1 \rangle\}$. It is immediate to check that $(\{G_\varphi, \bar{G}_\varphi\}, \top)$ is normalized—w.l.o.g., φ is not valid, i.e., there is an assignment that makes φ false. Moreover, note that $\text{EG}^*\text{-VAL}(\{G_\varphi, \bar{G}_\varphi\}, \top) = 1$ holds, if and only if, there is an interpretation σ where X evaluates true and such that $G_\varphi(\sigma) = 1$. Hence, $\text{EG}^*\text{-VAL}(\{G_\varphi, \bar{G}_\varphi\}, \top) = 1$ if, and only if, φ is satisfiable.

We conclude with the case $C = \{\wedge, \vee\}$ and $S = \{+, -\}$. Consider again the formula φ and let be $\{X_1, \dots, X_n\}$ the set of its variables. Let us build the formula $\bar{\varphi}$ where each occurrence of a literal $\neg X_i$ is replaced with the variable \bar{X}_i : $\bar{\varphi}$ is a formula over $\{X_1, \dots, X_n, \bar{X}_1, \dots, \bar{X}_n\}$ without negation, hence satisfiable. Let $G = \{\langle Z \wedge (\bigvee_i (X_i \wedge \bar{X}_i)), -1 \rangle, \langle \top, 1 \rangle\}$ and $\bar{G} = \{\langle Z, 1 \rangle\}$, where Z is a fresh variable. Note that the scenario $(\{G, \bar{G}\}, \bar{\varphi})$ is normalized, and that $\text{EG}^*\text{-VAL}(\{G, \bar{G}\}, \bar{\varphi}) = 1$ if, and only if, φ is satisfiable. \square

7 Conclusion

We have described and studied a general framework for group decision-making over combinatorial domains, where the representation of utility functions is made via weighted propositional formulas. We have characterized entirely the frontier of tractability with respect to meaningful restrictions on the number of agents, the number of formulas per agent, the number of different variables appearing in formulas, the syntax of formulas, and the signs of weights. Concerning the restrictions on the language, we point out that our fragments encode reasonable classes of preferences. For instance, *positive formulas* (without negation) are relevant for expressing preferences over sets of goods such as in fair division and auctions, while *cubes* (without disjunction) are relevant for representing the value of a coalition in coalitional games and also for some multiple referenda—expressive issues are analyzed in more details in [Uckelman *et al.*, 2009].

Our approach contrasts with the study of specific frameworks, such as those in [Escoffier *et al.*, 2013] and in other places. On the one hand, focusing on a specific problem (such as fair division) allows to obtain more precise, more specific results. On the other hand, the power of a general framework is that it allows to reason on new problems, before they are possibly investigated further. Avenues of further research include the study of other restrictions such as bounded weights, of the approximability of optimization for the intractable classes, and of determining the sets of agents who are most responsible for the loss of social welfare.

Acknowledgments

G. Greco's work was also supported by a Kurt Gödel Research Fellowship, awarded by the Kurt Gödel Society. J. Lang's work was also supported by the ANR-14-CE24-0007-01 project "CoCoRICo-CoDec".

References

- [Bansal and Sviridenko, 2006] N. Bansal and M. Sviridenko. The Santa Claus problem. In *Proc. of STOC'06*, pp. 31–40.
- [Betzler *et al.*, 2013] N. Betzler, A. Slinko, and J. Uhlmann. On the computation of fully proportional representation. *J. Artif. Intell. Res. (JAIR)*, 47:475–519, 2013.
- [Bezakova and Dani, 2005] I. Bezakova and V. Dani. Allocating indivisible goods. *SIGecom Exch.*, 5(3):1118, 2005.
- [Bordeaux *et al.*, 2014] L. Bordeaux, Y. Hamadi, and P. Kohli (editors). *Tractability: Practical Approaches to Hard Problems*. Cambridge University Press, 2014.
- [Boutilier and Hoos, 2001] C. Boutilier and H. Hoos. Bidding languages for combinatorial auctions. In *Proc. of IJCAI'01*, pp. 1211–1217.
- [Bouveret and Lang, 2008] S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *J. Artif. Intell. Res. (JAIR)*, 32:525–564, 2008.
- [Brams *et al.*, 2007] S. Brams, M. Kilgour, and R. Sanver. A minimax procedure for electing committees. *Public Choice*, 3-4(132):401–420, 2007.
- [Chamberlin and Courant, 1983] B. Chamberlin and P. Courant. Representative deliberations and representative decisions: Proportional representation and the borda rule. *A. Political Science Review*, 77(3):718–733, 1983.
- [Coste-Marquis *et al.*, 2004] S. Coste-Marquis, J. Lang, P. Liberatore, and P. Marquis. Expressive power and succinctness of propositional languages for preference representation. In *Proc. of KR'04*, pp. 203–212.
- [Elkind and Wooldridge, 2009] E. Elkind and M. Wooldridge. Hedonic coalition nets. In *Proc. of AAMAS'09*, pp. 417–424.
- [Elkind *et al.*, 2009] E. Elkind, L. Goldberg, P. Goldberg, and M. Wooldridge. A tractable and expressive class of marginal contribution nets and its applications. *Math. Log. Q.*, 55(4):362–376, 2009.
- [Elzinga *et al.*, 1976] J. Elzinga, D. Hearn, and W. Randolph. Minimax multifacility location with euclidean distances. *Transportation Science*, 10(4):321–336, 1976.
- [Endriss, 2011] U. Endriss. Logic and Social Choice Theory. *Logic and Philosophy Today*, Volume 2, Pages 333–377, College Publications, 2011.
- [Escoffier *et al.*, 2013] B. Escoffier, L. Gourvès, and J. Monnot. Fair solutions for some multiagent optimization problems. *Autonomous Agents and Multi-Agent Systems*, 26(2):184–201, 2013.
- [Gonzales *et al.*, 2008] C. Gonzales, P. Perny, and S. Queiroz. Preference aggregation with graphical utility models. In *Proc. of AAAI'08*, pp. 1037–1042.
- [Greco and Scarcello, 2013] G. Greco and F. Scarcello. Constraint satisfaction and fair multi-objective optimization problems: Foundations, complexity, and islands of tractability. In *Proc. of IJCAI'13*, pp. 545–551.
- [Jeong and Shoham, 2005] S. Jeong and Y. Shoham. Marginal contribution nets: A compact representation scheme for coalitional games. In *Proc. of EC'05*, pp. 193–202.
- [Lafage and Lang, 2000] C. Lafage and J. Lang. Logical representation of preferences for group decision making. In *Proc. of KR'00*, pp. 457–468.
- [Lang and Xia, 2015] J. Lang and L. Xia. Voting on Combinatorial Domains. In *Handbook on Computational Social Choice*, Chapter 9. Cambridge University Press, to appear.
- [Lu and Boutilier, 2011] T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proc. of IJCAI-11*, pp. 280–286.
- [Monroe, 1995] B. Monroe. Fully proportional representation. *A. Political Science Review*, 89(4):925–940, 1995.
- [Nisan, 2006] N. Nisan. Bidding languages for combinatorial auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial auctions*, chapter 9. MIT Press, 2006.
- [Papadimitriou, 2003] C. Papadimitriou. *Computational complexity*. John Wiley and Sons Ltd., 2003.
- [Procaccia *et al.*, 2008] A. Procaccia, J. Rosenschein, and A. Zohar. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 30(3):353–362, April 2008.
- [Shachnai and Tamir, 2001] H. Shachnai and T. Tamir. On two class-constrained versions of the multiple knapsack problem. *Algorithmica*, 29(3):442–467, 2001.
- [Skowron *et al.*, 2013] P. Skowron, P. Faliszewski, and A. Slinko. Achieving fully proportional representation is easy in practice. In *Proc. of AAMAS'13*, pp. 399–406.
- [Uckelman and Endriss, 2010] J. Uckelman and U. Endriss. Compactly representing utility functions using weighted goals and the max aggregator. *Artif. Intell.*, 174(15):1222–1246, 2010.
- [Uckelman *et al.*, 2009] J. Uckelman, Y. Chevaleyre, U. Endriss, and J. Lang. Representing utility functions via weighted goals. *Math. Log. Q.*, 55(4):341–361, 2009.
- [van Hentenryck and Deville, 1991] P. van Hentenryck and Y. Deville. The cardinality operator: A new logical connective for constraint logic programming. In *Proc. of ICLP'91*, pp. 745–759.
- [Wagner, 1987] K. Wagner. More complicated questions about maxima and minima, and some closures of NP. *Theoretical Computer Science*, 51(1):53–80, 1987.