

A General Approach for Computing a Consensus in Group Decision Making That Integrates Multiple Ethical Principles

Francisco Salas-Molina^{*1}, Filippo Bistaffa², and Juan A.
Rodríguez-Aguilar²

¹*Universitat Politècnica de València, Ferrandiz y Carbonell s/n, 03801 Alcoy, Spain*

²*IIIA-CSIC, Campus UAB, 08913 Bellaterra, Spain*

January 15, 2024

Abstract

We tackle the problem of computing a consensus according to multiple *ethical principles*—which can include, for example, the principle of maximum freedom associated with the Benthamite doctrine and the principle of maximum fairness associated with the Rawlsian principles— among the preferences of different individuals in the context of *Group Decision-Making* (GDM). More formally, we put forward a novel formalisation of the above-mentioned problem based on a *multi- ℓ_p -norm approximation problem* that aims at minimising multiple p -metric distance functions, where each parameter p represents a given ethical principle. Our contribution incurs obvious benefits from a social-choice perspective. Firstly, our approach significantly generalises state-of-the-art approaches that were limited to only two ethical principles ($p = 1$, for maximum freedom, and $p = \infty$, for maximum fairness). Secondly, our experimental results considering an established test case demonstrate that our approach is capable, thanks to a novel re-weighting scheme, to compute a multi-norm consensus that takes into account each ethical principle in a balanced way, in contrast with state-of-the-art approaches that were heavily biased towards the $p = 1$ ethical principle.

Keywords: reference aggregation, distance functions, consensus space.

1 Introduction

Group Decision-Making (GDM) is ubiquitous in everyday life. Even when decision-makers act individually, they often receive advice or suggestions from others. Thus, decisions are often social in nature and involve multiple group members (Tindale and Winget, 2019; Yazdani et al., 2022). Among others, one

^{*}Corresponding author. E-mail addresses: frasamo@upv.es

prominent example of GDM is *citizen participation* (Fagence, 2014), adopted by the public administrations of many major cities all over the world. As an example, the citizen participation platform by the Barcelona administration¹ is actively being used to decide and implement local policies with the direct involvement of citizens. Along these lines, reaching an agreement that represents a *principled* compromise among multiple preferences or opinions is of utmost importance in modern society.

Over the years, a vast wealth of literature concerned with GDM has formulated and solved this problem in many different ways (Butler and Rothstein, 2006; Xu and Cai, 2011; Herrera-Viedma et al., 2014; Ortega Riejos et al., 2015; Pérez et al., 2018; Zhang et al., 2020; Lu et al., 2021; Noori et al., 2021; Pamucar et al., 2022; Zhang and Chen, 2022; Liu et al., 2022; Liang et al., 2022; Hassani et al., 2022; Zhang et al., 2021; Xiao et al., 2022). In this paper, we follow the approach introduced by González-Pachón and Romero (1999, 2008, 2011, 2016), where the problem of finding a consensus or a collective decision is achieved by minimising a p -metric distance function among the preferences (or opinions) of the individuals.

One major benefit of this approach is that it allows decision-makers to compute a consensus following different *ethical principles* according to different values of the parameter p . For instance, $p = 1$ corresponds to the idea of maximum freedom of the individuals derived from the theory of utilitarianism by Bentham (1789), whereas $p = \infty$ corresponds to the principle of maximum fairness, from the Rawlsian idea of considering only the welfare of the worst-off group (Rawls, 1973).² The principle of maximum freedom is relevant when the increase in the welfare of a well-off group by one unit has the same social value as an increase in the welfare of a disadvantaged group by one unit. On the contrary, in the principle of maximum fairness, the welfare of society only depends on the utility of the worst-off individual or social group. For instance, let us consider the problem faced by policy-makers in deciding whether to construct a new parking lot or a bike lane. Assuming that the majority of the population use cars as a personal means of transportation, according to the maximum freedom principle the construction of a new parking lot yields the maximum welfare for the society. On the other hand, the worst-off group (i.e., cyclists) would benefit from the construction of a new bike lane, hence it would adhere to the maximum fairness principle.

Apart from these two well-known cases, an exciting research question worthy of investigation is how to compute a consensus for any of the infinite values of $p \in \mathbb{N}$ that represent *intermediate* ethical principles other than maximum freedom or maximum fairness. Unfortunately, the approach proposed by González-Pachón and Romero (2016) cannot be applied to values of p different from 1 and ∞ , due to the computational limitations of the goal programming solution technique adopted by the authors.

Besides that, and taking a more general perspective, computing a consensus may require the involvement of *multiple* ethical principles, be it because the decision-making involves multiple stakeholders (each one with their own ethical principle), or because the consensus aims at a trade-off between ethical principles. To the best of our knowledge, despite previous attempts considering ad-hoc

¹ Accessible at <https://ajuntament.barcelona.cat/participaciociudadana>.

² In this paper, we use “norm”, “metric” and “ethical principle” to refer to the same concept.

combinations of $p = 1$ and $p = \infty$ (González-Pachón and Romero, 2016), computing a consensus that involves a general combination of multiple, any-valued ethical principles is still an open problem that has not been addressed by the literature.

Against this background, in this paper we propose a novel generalisation to the previous approach by González-Pachón and Romero (1999, 2008, 2011, 2016) based on an ℓ_p -norm approximation problem (Boyd and Vandenberghe, 2004) (also called ℓ_p -regression (Adil et al., 2019)), which is able to account for multiple, any-valued ethical principles. Thus, we show that computing a consensus for a multi-principle problem amounts to solving a multi- ℓ_p -norm approximation.

Our generalisation is conceived to provide expressiveness for: (i) each stakeholder to choose their own ethical principle; and (ii) setting the relative importance of ethical principles, and hence the preferences over ethical principles (e.g., three different groups supporting three ethical principles setting the relative importance by means of the size of the group). Such a novel approach yields multiple benefits both from the theoretical and the computational perspectives. Most importantly, our approach can be used to compute a consensus for a given set of ethical principles, and also in particular for any value of p (not only for $p = 1$ and $p = \infty$), hence representing a significant generalisation of state-of-the-art work on the matter.

Importantly, our approach has been conceived to pursue a *balanced* contribution of each ethical principle to the consensus. This is achieved by carefully aggregating the multiple p -norms involved in the computation of a consensus. More precisely, we introduce a *re-weighting approach* for multi- ℓ_p -norm approximation that aims at *normalising* the values of the different p -norms involved prior to their aggregation. Finally, following the methodology proposed in González-Pachón and Romero (2016), we empirically confirm that our re-weighting approach is indeed necessary to fairly include the contributions from the different ethical principles at hand.

In more detail, this paper advances the state-of-the-art in the following ways:

1. We show that the problem of computing the consensus among individuals under multiple p -metric distance functions (González-Pachón and Romero, 2016) can be cast as a *multi-norm approximation problem*, a particular and novel type of norm approximation problem (Boyd and Vandenberghe, 2004). We discuss how this enables the use of convex optimisation solution techniques as the means to compute the consensus for *multiple* ethical principles.
2. By doing so, our multi- ℓ_p -norm approximation approach provides a general theoretical framework that can accommodate *many* values of p (i.e., many ethical principles), hence generalising the approach in González-Pachón and Romero (2016). Furthermore, our approach also allows us to compute the consensus for *any* single value of p (in contrast with González-Pachón and Romero (2016), which can only deal with $p = 1$ and $p = \infty$).
3. We also propose a *re-weighting approach* for our multi- ℓ_p -norm approximation (henceforth referred to as *weighted model*) so that the ethical principles involved fairly contribute to the consensus, improving upon the approach proposed in González-Pachón and Romero (2016). This is achieved

by *normalising* the values of the different p -norms involved in computing the consensus.

4. Considering an established test case, we illustrate our general methodology for single ethical principles (besides $p = 1$ and $p = \infty$) by analysing the distribution of the differences in individual preferences with respect to the consensus for several values of p . We exploit this new capability provided our general approach to empirically analyse the consensus for the whole space of intermediate values of p , highlighting that values different from $p = 1$, $p = 2$ and $p = \infty$ may have a complex semantic interpretation.
5. Considering the same test case, we empirically analyse how considering different combinations of multiple ethical principles in our multi- ℓ_p -norm approximation model impacts the consensus, both for the unweighted model (without re-weighting) and for the weighted one (with re-weighting). We observe that using an unweighted model would lead to a biased consensus since the contribution of the utilitarian principle ($p = 1$) dominates the contributions of other ethical principles. Nonetheless, we also observe that our weighted model manages to correct such bias. Thus, this model leads to a balanced contribution of ethical principles to the consensus.

The structure of this paper is as follows. Section 2 reviews related works. Section 3 provides useful background about the use of distance functions in social choice theory. Section 4 describes our novel ℓ_p -norm approximation approach to account for multiple ethical principles,

the so-called *multi- ℓ_p -norm approximation*. Section 5 introduces a *re-weighting approach* for multi- ℓ_p -norm approximation that aims at ensuring the balanced contribution of multiple p -norms to the consensus. Section 6 empirically analyses the semantics of single ethical principles and analyses the impact of considering multiple ethical principles on a consensus. Section 7 offers some concluding remarks and a natural extension of this work.

2 Literature review

The problem of computing a consensus or an aggregation of preferences of different individuals (or within social groups) has been widely studied in the literature, which usually considers individuals that have to express their opinions on some alternatives or policies (Butler and Rothstein, 2006). In many cases, preferences over alternatives are expressed by preference relations (Xu and Cai, 2011). In this context, authors have considered a variety of different settings, including dynamic scenarios (Pérez et al., 2018; Hassani et al., 2022), scenarios with fuzzy preferences (Herrera-Viedma et al., 2014; Zhang et al., 2020; Zhang and Chen, 2022) or characterised by a social network structure (Lu et al., 2021; Liu et al., 2022; Liang et al., 2022). Some authors study the impact of the social trust relationships on assessments-modifications in the consensus reaching in social network group decision making Zhang et al. (2021). Other authors describe a multi-stage consensus optimisation model with bounded confidence to obtain the adjustment suggestions by maximising the level of consensus Xiao et al. (2022). If the predetermined level of consensus cannot be reached, the adjustment suggestions obtained by the model are adopted to guide the preference modification of the group members.

On the one hand, a body of literature has focused on the specific case where each individual expresses their preferences in the form of a *ranking*. Finding this optimal aggregated ranking can be formulated as an optimisation problem, which is usually NP-hard (Dwork et al., 2001). Thus, several heuristic *ranking aggregation methods* have been proposed, including batch mode methods, instant-runoff mode methods, and, more recently, hierarchical methods (Ding et al., 2018; Aledo et al., 2021). Other related works propose methods to overcome the difficulties of eliciting preferences over too many alternatives including a branch-and-bound algorithm to construct a consensus ranking (Cook et al., 2010), an efficient method for aggregating measurements acquired by an uncalibrated environmental sensory network (Fishbain and Moreno-Centeno, 2016), an axiomatic method to aggregate a set of incomplete rankings into a consensus ranking (Moreno-Centeno and Escobedo, 2016), a binary programming formulation for the generalised Kemeny (Kemeny and Snell, 1962) rank aggregation problem (Yoo and Escobedo, 2021), and a multimodal data aggregation methodology for jointly aggregating heterogeneous ordinal and cardinal evaluation inputs into a consensus evaluation (Escobedo et al., 2022).

On the other hand, in some situations it might be difficult or even impossible for individuals to express their preferences as rankings, hence the above-discussed methods cannot be applied. Indeed, in this paper, we follow an alternative and more general approach proposed in another strand of literature (González-Pachón and Romero, 1999, 2008, 2011, 2016) where preferences are expressed by a *pairwise representation*, i.e., a square matrix R where the value $R_{i,j}$ represents the preference of object i with respect to object j (see Section 3 for more details). However, we should mention that pairwise comparison matrices require m^2 inputs from each individual whereas a ranking requires only m inputs.

In this context, the problem of finding a consensus between individuals (or within social groups) can be approached by means of a p -metric distance function that allows consideration of several ethical principles by varying the value of $p \in \mathbb{R}$ belonging to the closed interval $[1, \infty]$ as proposed by González-Pachón and Romero (1999, 2008, 2011, 2016). Specifically, González-Pachón and Romero (2011, 2016) consider particular values of p , evaluating the quality of consensus among groups by the aggregated value of the distance between the position of each group and the consensus point. To solve the problem, the authors relied on extended goal programming (Romero, 2001; Jones et al., 2010), which can only be used in the case of particular values of p (i.e., $p = 1$ and $p = \infty$), as using such a technique for $p > 1$ is computationally too demanding.

Finally, in the context of measurement of economic inequality, (Atkinson et al., 1970) proposed the use of a parametric inequality measure in which a parameter in the range $[1, \infty]$ plays the role of a degree of inequality aversion or the relative sensitivity to transfers at different income levels. Instead, in this paper we use p as a degree of utilitarian aversion.

In what follows, we provide detailed background on the social choice problem based on parametric distance functions following the notation used in González-Pachón and Romero (2016).

3 Parametric distance functions in social choice theory

In this section, we describe the social choice problem based on p -metric distance functions following the notation used in González-Pachón and Romero (2016). We consider a society with n individuals indexed by $i = 1, 2, \dots, n$ that has to provide judgements on m objects or alternatives indexed by $j, k = 1, 2, \dots, m$. These judgements can be cardinal, when the individuals express some degree of preference over the alternatives, and ordinal when they use a Boolean value to express their preference for one alternative over another one. The rest of the elements of the problem are as follows:

- Weight w_i represents the social influence of the i th individual (or social group).
- Judgement R_{jk}^i is the value provided by the i th individual when comparing the j th and the k th object.
- Judgement R_{jk}^S is the value assigned by the society as a whole when comparing the j th and the k th object. These are the unknowns, the solution that we seek.
- Set \mathbf{F} includes the constraints that solutions R_{jk}^S must satisfy.

From the previous definitions, the weighted Minkowski p -metric distance function (U_p) can be used as a generator of social choice functions as described in González-Pachón and Romero (1999, 2008, 2011, 2016). The ultimate goal is obtaining social consensus points R_{jk}^S such that the weighted p -metric distance between R_{jk}^S and judgement values R_{jk}^i provided by the individuals within the society is minimised. Here, p is a real number in the closed interval $[1, \infty]$, which is a key parameter in the following distance function:

$$U_p = - \left[\sum_{i=1}^n \sum_{j=1}^m \sum_{k=1, k \neq j}^m w_i^p |R_{jk}^i - R_{jk}^S|^p \right]^{1/p}. \quad (1)$$

From this distance function, González-Pachón and Romero (2016) derived three particular cases to represent well-known ethical principles. By setting $p = 1$, the utilitarian principle of freedom to maximise the total welfare U_B proposed by Bentham (1789) is represented as:

$$U_B = - \left[\sum_{i=1}^n \sum_{j=1}^m \sum_{k=1, k \neq j}^m w_i |R_{jk}^i - R_{jk}^S| \right]. \quad (2)$$

By setting $p = \infty$, the idea of fairness is represented by the min-max principle proposed by Rawls (1973). This principle, denoted by U_R , implies that the maximum deviation from an individual judgement R_{jk}^i with respect to the consensus point R_{jk}^S is minimised:

$$U_R = -\max [w_i |R_{jk}^i - R_{jk}^S|]. \quad (3)$$

Finally, the idea of maximum equity by Marx is derived from U_B in Equation (2), when considering the following constraint:

$$w_1|R_{jk}^1 - R_{jk}^S| = w_2|R_{jk}^2 - R_{jk}^S| = \dots = w_n|R_{jk}^n - R_{jk}^S|. \quad (4)$$

To address the computational problems of the above models, González-Pachón and Romero (2016) proposed the following change in variables derived from an initial proposal by Charnes and Cooper (1977):

$$n_{jk}^i = \frac{1}{2} [|R_{jk}^i - R_{jk}^S| + (R_{jk}^i - R_{jk}^S)] \quad (5)$$

$$p_{jk}^i = \frac{1}{2} [|R_{jk}^i - R_{jk}^S| - (R_{jk}^i - R_{jk}^S)]. \quad (6)$$

Then, by adding Equations (5) and (6), and by subtracting Equation (6) from Equation (5), the following identities are obtained:

$$n_{jk}^i + p_{jk}^i = |R_{jk}^i - R_{jk}^S| \quad (7)$$

$$n_{jk}^i - p_{jk}^i = R_{jk}^i - R_{jk}^S. \quad (8)$$

By using Equation (7) and Equation (8), objective function in Equation (1) becomes:

$$U_p = - \left[\sum_{i=1}^n \sum_{j=1}^m \sum_{k=1, k \neq j}^m w_i^p (n_{jk}^i - p_{jk}^i)^p \right]^{1/p} \quad (9)$$

subject to Equation (8) and $R_{jk}^S \in \mathbf{F}$. As a result, the Benthamite solution $U_B = U_1$ is obtained when $p = 1$. To obtain the Rawlsian solution U_R , the following goal programming Chebyshev model is solved to obtain D , which represents the disagreement of the member of the society with the opinions most displaced from the solution obtained:

$$U_R = -D \quad (10)$$

subject to Equation (8), $R_{jk}^S \in \mathbf{F}$, and:

$$w_i \sum_{j=1}^m \sum_{k=1, k \neq j}^m (n_{jk}^i - p_{jk}^i) - D \leq 0. \quad (11)$$

Finally, although it is likely that the system of equations in Equation (4) has no solution, the idea of maximum equity (U_M) can be approximated by relying on meta-goal programming, as described in González-Pachón and Romero (2016) and Caballero et al. (2006):

$$U_M = - \sum_{i=1}^{n-1} \sum_{t=i+1}^n (\eta_{it} + \rho_{it}) \quad (12)$$

subject to:

$$\begin{aligned} w_i(n_{jk}^i + p_{jk}^i) - w_t(n_{jk}^t + p_{jk}^t) + \eta_{it} - \rho_{it} &= 0 \quad i = 1, \dots, n-1 \quad t = i+1, \dots, n \\ R_{jk}^S + n_{jk}^t - p_{jk}^t &= R_{jk}^i \quad i = 1, \dots, n \quad \forall j, k \\ R_{jk}^S &\in \mathbf{F}. \end{aligned}$$

where η_{it} and ρ_{it} are the usual deviation variables used in goal programming.

From these particular cases, a convex combination of U_B , U_R and U_M can be derived to account for three different ethical principles by varying control parameters λ_1 and λ_2 in the interval $[0, 1]$:

$$U(\lambda_1, \lambda_2) = \lambda_1 U_B + \lambda_2 U_R + (1 - \lambda_1 - \lambda_2) U_M, \quad (13)$$

subject to the set of restrictions in \mathbf{F} and $\lambda_1 + \lambda_2 \in [0, 1]$. Note that for $\lambda_1 = 1$ and $\lambda_2 = 0$, we obtain the Benthamite solution of maximum freedom; for $\lambda_1 = 0$ and $\lambda_2 = 1$, we obtain the Rawlsian solution of maximum fairness; and for $\lambda_1 = 0$ and $\lambda_2 = 0$, we obtain the Marxian solution of maximum equity.

As for the set \mathbf{F} of constraints that the solution R_{jk}^S must satisfy, three sets of restrictions were considered by González-Pachón and Romero (2008), allowing the solution to reflect any valid set of cardinal and/or ordinal preferences:

1. Information about preferences is ordinal and complete. In this case, alternatives are ordered from the best to the worst by using a linear order such that $r_{jk}^i = 1$ and $r_{kj}^i = 0$ if the i th decision-maker prefers the j th alternative to the k th alternative, and $r_{jk}^i = 0$ and $r_{kj}^i = 1$ if the i th decision-maker prefers the k th alternative to the j th alternative. This assumption implies that $r_{jk}^i + r_{kj}^i = 1, \forall i, j, k$ and $a_i^j = \sum_{k=1}^m r_{jk}^i$ provides a valuation R_j^i for the j th alternative. Then, Borda's count is used to define feasible set \mathbf{F} as follows:

$$\mathbf{F} = \left\{ R_j^S \mid R_j^S \in \mathbb{R}^+, 1 \leq R_j^S \leq m, \sum_{j=1}^m R_j^S = \frac{(m+1)m}{2} \right\} \quad (14)$$

where, by convention, the number m is assigned to the best alternative and number 1 to the worst. This definition is relevant as it is the key reason that \mathbf{F} represents a convex set, hence ensuring feasibility in linear programming in direct contrast to other common social choice measures on rankings such as Kendall tau, Hamming and Kemeny distance functions.

2. Information about preferences is ordinal and partial. In this case, some alternatives may be incomparable and this assumption implies that $r_{jk}^i + r_{kj}^i \leq 1 \forall i, j, k$. Now, a_i^j , defined as in the first case, establishes a minimum bound for R_j^i , and $b_i^j = m - \sum_{k=1}^m r_{jk}^i$ establishes a maximum bound for R_j^i . Then $R_j^i \in [a_i^j, b_i^j]$ and feasible set \mathbf{F} is defined as in Equation (14).

3. Information about preferences is cardinal and complete. In this particular case, a valued binary preference relation reports some degree of the cardinality of preferences from one alternative over another. Decision-makers' valuations should belong to a certain closed interval $[t_1, t_2]$ as it was proposed by Saaty (1977), where it is also assumed that $R_{jk}^i = 1/R_{kj}^i$. These assumptions lead to the following set \mathbf{F} :

$$\mathbf{F} = \left\{ R_{jk}^S \mid t_1 \leq R_{jk}^S \leq t_2 \right\}. \quad (15)$$

From the observation of Equation (1) and Equation (13), a straightforward research question arises. If three different ethical principles can be considered by optimising a convex combination of different p -metric distance functions subject to some set of constraints, one may be able to explore a larger space of multiple ethical principles in a diverse society if additional p -metric distances are considered. In what follows, we further elaborate on this research question.

4 A multi- ℓ_p -norm approximation approach

In this section, we propose a novel generalisation to the previous approach by González-Pachón and Romero (1999, 2008, 2011, 2016) based on an ℓ_p -norm approximation problem, which is able to account for multiple, any-valued ethical principles.

Along the lines of Equation (13), we now provide a formal definition of the multi-principle problem. Our definition is conceived to provide expressiveness for: (i) each stakeholder to choose their own ethical principle; and (ii) setting the relative importance of ethical principles, and hence the preferences over ethical principles (e.g., three different groups supporting three ethical principles setting the relative importance by means of the size of the group).

Definition 1. We consider a non-negative weighted sum of p -metric distances, each representing the corresponding U_p objective function. Specifically, we consider a set \mathbf{P} of multiple ethical principles and a set $\boldsymbol{\lambda} \subseteq \mathbb{R}_{\geq 0}$, where $\lambda_p \in \boldsymbol{\lambda}$ represents the non-negative weight associated with ethical principle p .

We define the problem of computing the consensus R^S considering the above-defined sum of ethical principles as the problem of computing

$$R^S = \arg \min_R \sum_{p \in \mathbf{P}} \lambda_p \left[\sum_{i=1}^n \sum_{j=1}^m \sum_{k=1, k \neq j}^m w_i^p |R_{jk}^i - R_{jk}|^p \right]^{1/p} \quad (16)$$

subject to the set of restrictions in \mathbf{F} .

We now show that solving Equation (16) is equivalent to solving a *multi- ℓ_p -norm approximation problem*. More specifically, we begin by defining the *vectorisation* operation, which we will employ in Lemma 1 to show that solving Equation (16) for a single ethical principle is equivalent to solving a standard ℓ_p -norm approximation problem (Boyd and Vandenberghe, 2004). Then, in Theorem 1 we leverage Lemma 1 to show that solving Equation (16) is equivalent to solving a *multi- ℓ_p -norm approximation problem*.

Definition 2. Given an $n \times m$ matrix M , the vectorisation operation $\text{vec}(M)$ produces a vector v of $n \cdot m$ elements obtained by arranging the elements of M in row-major order, i.e., by arranging them sequentially row by row. Formally, $v_{i \cdot m + j} = M_{ij} \forall i \in \{1, \dots, n\} \forall j \in \{1, \dots, m\}$.

Lemma 1. The problem of computing the consensus

$$R^S = \arg \min_R \left[\sum_{i=1}^n \sum_{j=1}^m \sum_{k=1, k \neq j}^m w_i^p |R_{jk}^i - R_{jk}|^p \right]^{1/p} \quad (17)$$

subject to \mathbf{F} is equivalent to solving the ℓ_p -norm approximation problem

$$\begin{aligned} & \text{minimise } \overbrace{\|Ax - b\|_p}^{\eta_p}, \\ & \text{subject to } \mathbf{F}, \end{aligned} \quad (18)$$

where

$$A = \begin{bmatrix} w_1 \cdot I_{m^2} \\ \vdots \\ w_n \cdot I_{m^2} \end{bmatrix}, \quad b = \begin{bmatrix} w_1 \cdot \text{vec}(R^1) \\ \vdots \\ w_n \cdot \text{vec}(R^n) \end{bmatrix},$$

and $I_{m^2} \in \mathbb{R}^{m^2 \times m^2}$ is the identity matrix of size m^2 . For convenience, we denote as η_p the value of the objective function of Equation (18), which will be employed later in Equation (23).

Proof. Without loss of generality, we assume that $R_{jk} = 0$ when $j = k$. As a consequence, Equation (17) is equivalent to

$$\arg \min_{R^S} \left[\sum_{i=1}^n \sum_{j=1}^{m^2} |w_i \cdot \text{vec}(R^S)_j - w_i \cdot \text{vec}(R^i)_j|^p \right]^{1/p}. \quad (19)$$

We can further rewrite (19) as

$$\arg \min_{R^S} \left[\sum_{i=1}^n \|w_i \cdot \text{vec}(R^S) - w_i \cdot \text{vec}(R^i)\|^p \right]^{1/p}. \quad (20)$$

We define $A \in \mathbb{R}^{n \cdot m^2 \times m^2}$ and $b \in \mathbb{R}^{n \cdot m^2}$ as

$$A = \begin{bmatrix} w_1 \cdot I_{m^2} \\ \vdots \\ w_n \cdot I_{m^2} \end{bmatrix}, \quad b = \begin{bmatrix} w_1 \cdot \text{vec}(R^1) \\ \vdots \\ w_n \cdot \text{vec}(R^n) \end{bmatrix},$$

where $I_{m^2} \in \mathbb{R}^{m^2 \times m^2}$ is the identity matrix of size m^2 .

By exploiting the above-defined A and b , we can now formulate Equation (20) as an ℓ_p -norm approximation problem, i.e.,

$$\begin{aligned} & \text{minimise } \|Ax - b\|_p, \\ & \text{subject to } \mathbf{F}. \end{aligned}$$

Notice that the solution to the ℓ_p -norm approximation problem, i.e., the vector x , is the vectorisation of R^S subject to \mathbf{F} . \square

Now we are ready to prove that computing the consensus for a set of ethical principals, as defined by Equation (16), amounts to solving a multi- ℓ_p -norm approximation problem.

Theorem 1. *We consider a set \mathbf{P} of multiple of ethical principles and a set $\boldsymbol{\lambda} \subseteq \mathbb{R}_{\geq 0}$, where $\lambda_p \in \boldsymbol{\lambda}$ represents the non-negative weight associated with ethical principle p . The problem of computing the consensus*

$$R^S = \arg \min_R \sum_{p \in \mathbf{P}} \lambda_p \left[\sum_{i=1}^n \sum_{j=1}^m \sum_{k=1, k \neq j}^m w_i^p |R_{jk}^i - R_{jk}|^p \right]^{1/p}$$

subject to the set of restrictions in \mathbf{F} is equivalent to solving the multi- ℓ_p -norm approximation problem

$$\begin{aligned} & \text{minimise} \quad \sum_{p \in \mathbf{P}} \lambda_p \|Ax - b\|_p, \\ & \text{subject to} \quad \mathbf{F}, \end{aligned} \tag{21}$$

where

$$A = \begin{bmatrix} w_1 \cdot I_{m^2} \\ \vdots \\ w_n \cdot I_{m^2} \end{bmatrix}, \quad b = \begin{bmatrix} w_1 \cdot \text{vec}(R^1) \\ \vdots \\ w_n \cdot \text{vec}(R^n) \end{bmatrix}.$$

and $I_{m^2} \in \mathbb{R}^{m^2 \times m^2}$ is the identity matrix of size m^2 .

Proof. By direct application of Lemma 1 to each element of the sum. \square

Notice that, while the concept of ℓ_p -norm approximation problem has been deeply studied in the optimisation literature (Boyd and Vandenberghe, 2004), its generalisation to the multi- ℓ_p -norm case seems to have received very little attention. To the best of our knowledge, the only work attempting to formalise and solve an ℓ_p -norm approximation problem involving multiple norms is the one by Samejima and Matsushita (2016), even though the approach proposed by the authors is not suitable to solve Equation (16). Firstly, the *general norm approximation problem* proposed in Samejima and Matsushita (2016) is not exactly equivalent to Equation (16), since it involves the minimisation of a sum of *power of norms*, rather than considering norms directly. This is a formally different problem that is not aligned with the one proposed here, nor with previous literature (González-Pachón and Romero, 1999, 2008, 2011, 2016). Secondly, the algorithm proposed in Samejima and Matsushita (2016) is based on a numerical technique called *Iteratively Reweighted Least Squares* (IRLS), which is not guaranteed to converge for $p \geq 3$. Finally, such an algorithm can only be used to solve *unconstrained* problems, i.e., it would disregard the restrictions in the set \mathbf{F} . These limitations further motivate the need of a general and reliable method to solve the consensus computation problem involving multiple ethical principles, such as the one presented here.

Thanks to Theorem 1 we can solve Equation (16) by formulating it as in Equation (21), hence enabling the use of Convex Optimisation solution techniques.

Remark 1. *If the restrictions in the set \mathbf{F} are convex, the problem in Equation (18) is also convex, since it is well known that norms are convex functions (Boyd and Vandenberghe, 2004, Section 3.1.5). Similarly, the problem in Equation (21) is also convex, since it is a non-negative weighted sum of convex functions (Boyd and Vandenberghe, 2004, Section 3.2.1).*

Specifically, we solve Equation (21) by employing the techniques discussed by Alizadeh and Goldfarb (2001), Section 2.3g, who show how to represent inequalities involving p -norms as *Second-Order Cone Programming* constraints, enabling the use of off-the-shelf solvers such as CPLEX.

Remark 2. *The set of restrictions \mathbf{F} in Equation (16) (or, equivalently, in Equation (21)) can be used to impose certain properties on the resulting consensus R^S , for instance to obtain a complete total ranking (González-Pachón*

and Romero, 2008) among the m available options. Imposing such restrictions inherently reduces the set of valid solutions, which may result in a consensus R^S corresponding to a lower utility (as defined in Equation (1)) compared to the unrestricted case.

5 A re-weighting approach for multi- ℓ_p -norm approximation

The main purpose of Equation (21), as well as of the original Equation (13) by González-Pachón and Romero (2016), is to compute a consensus R^S that is the result of the aggregation of multiple ethical principles, or multiple p -norms. Unfortunately, this approach alone does not always yield a result with the above-mentioned semantic, as illustrated by the following example.

Let us consider the example in which one aims at computing the multi-norm consensus for $\lambda_1 = \lambda_2 = 0.5$ in Equation (13) or, equivalently, $\mathbf{P} = \{1, \infty\}$ each weighted with the same λ_p in Equation (21). Intuitively, the expected result should be a balanced mix of the single components (i.e., $p = 1$ and $p = \infty$), since both are given the same importance (i.e., the same λ).

Nonetheless, directly summing U_B and U_R in Equation (13) disregards the fact that U_B has a much higher value with respect to U_R (i.e., 184.17 *vs* 4 in the example considered in our experiments) and hence, it receives a much higher importance in the optimisation. Indeed, in this example the multi-norm consensus computed for $\mathbf{P} = \{1, \infty\}$ is *exactly equal* to the one computed for the single-norm case with $p = 1$. In other words, the inclusion of $p = \infty$ has not a measurable impact in the resulting consensus, which is, of course, in contrast with the expected semantic. More in general, this issue is due to the well known fact that different p -norms (or, equivalently, different U_p functions) have different scales depending on p , which has no effect for the single-norm case, but must be taken into account for the multi-norm case.

To tackle this issue, we propose a *re-weighting approach* for multi- ℓ_p -norm approximation that aims at *normalising* the values of the different p -norms. Let us first recall our definition of multi- ℓ_p -norm approximation in Equation (21):

$$\begin{aligned} & \text{minimise } \sum_{p \in \mathbf{P}} \lambda_p \|Ax - b\|_p, \\ & \text{subject to } \mathbf{F}. \end{aligned}$$

Our re-weighting approach seeks to determine the weight $\lambda_p \in \boldsymbol{\lambda}$ for each ethical principle $p \in \mathbf{P}$ with the goal of minimising $\Psi(\mathbf{P}, \boldsymbol{\lambda})$, i.e., the *variance* of the set of values corresponding to each component $\lambda_p \|Ax - b\|_p$, as defined hereafter.

Definition 3. *Given a set of ethical principles \mathbf{P} and vector of weights $\boldsymbol{\lambda}$, we define the measure*

$$\Psi(\mathbf{P}, \boldsymbol{\lambda}) = \text{Var}(\{\lambda_p \|Ax - b\|_p \mid p \in \mathbf{P}, \lambda_p \in \boldsymbol{\lambda}\}), \quad (22)$$

which quantifies how unbalanced is the sum of the components corresponding to each ethical principle $p \in \mathbf{P}$.

With the goal of minimising the above-defined Ψ measure, we propose the following re-weighting approach.

Definition 4. We define the re-weighted multi- ℓ_p -norm approximation problem as

$$\begin{aligned} & \text{minimise } \sum_{p \in \mathbf{P}} \frac{1}{\eta_p} \cdot \|Ax - b\|_p, \\ & \text{subject to } \mathbf{F}, \end{aligned} \tag{23}$$

where η_p is the value of the objective function of the single ℓ_p -norm approximation problem in Equation (18) considering A and b (see previous definition in Lemma 1).

In Equation (23) each term of the multi-norm sum is divided for the value of $\|\cdot\|_p$ in the single-norm case, so as to normalise its contribution in the sum and, as a consequence, minimise the measure Ψ that we defined as our objective. Notice that computing each η_p requires solving each single-norm approximation problem in Equation (18) as a preliminary step to then solve Equation (23).

Remark 3. Following Remark 1, it is easy to show that Equation (23) is also a convex optimisation problem if the restrictions in the set \mathbf{F} are convex.

6 Experimental results

In this section, we experimentally evaluate our approach following the methodology previously proposed in the literature (González-Pachón and Romero, 2016). Specifically, we consider the participatory forest plan problem (Nordström et al., 2009) carried out in Lycksele (Sweden), where several stakeholders were questioned about their preferences with respect to several criteria. Seven recreationists were interviewed about the relative importance attached to five criteria, following a pairwise comparison format. The original 5×5 pairwise comparison matrices can be found in González-Pachón and Romero (2016).

Given the above experimental scenario, our objectives are the following:

1. Analyse the distribution of the differences of the individual preferences with respect to the consensus (i.e., the residuals of the ℓ_p -norm approximation problem) computed by using our model in Equation (18) for several values of p .
2. Analyse how considering different combinations of multiple ethical principles in our multi- ℓ_p -norm approximation model impacts such a distribution of differences, both for the unweighted model (Equation (21)) and for the weighted one (Equation (23)).³

6.1 Single ethical principle case

One of the main contributions of our work is generalising previous consensus computation approaches (González-Pachón and Romero, 2016) by providing the opportunity of considering *any* ethical principle p (and not only $p \in \{1, \infty\}$). Along these lines, in our first experiment we compute the consensus by using our model in Equation (18) for $p \in \{1, 2, 3, 4, 5, 10, 50, 100, 500, \infty\}$. Then, in Figure

³Our code is available at <https://github.com/filippobistaffa/social-choice-pnorm>.




Figure 1: Boxplots representing the distribution of residuals for different ethical principles.

1 we show the boxplots of the distribution of the differences of the individual preferences with respect to the consensus (i.e., the residuals of the corresponding ℓ_p -norm approximation problem) computed for each p . Here we report such boxplots (rather than reporting the actual consensus matrix R^S) because they allow us to compare the distributions of the residuals in a visual way for a large number of ethical principles, which is the goal of this experiment.

As expected, $p = 1$ (i.e., fully utilitarian case) is the ethical principle that minimises the *average* difference ($\mu_1 \simeq 1.05$) among all possible choices, since considering the ℓ_1 -norm amounts to minimising the *total* sum of differences and, hence, the average one as well. The ℓ_1 -norm is also associated with the highest *maximum* difference ($\max_1 = 6$).

On the other hand, $p = 2$ results in the consensus that minimises the *mean squared error*, i.e., it results in the distribution of differences with the lowest *variance* ($\sigma_2^2 = 1.45$, see Table 1), corresponding to the boxplot with the lowest *interquartile range* among all ethical principles.

Intuitively, we can interpret this result as follows: for $p = 2$ the individual differences are more compactly distributed, i.e., they are less sparsely distributed. More in general, $p = 2$ can be interpreted as the consensus that aims at the most *homogeneous* difference across all individuals, rather than simply minimising the maximum one (as done for $p = \infty$).

Moreover, we immediately observe that increasing p (i.e., shifting to a more egalitarian consensus) results in a *lower* maximum difference ($\max_2 = 4.57$) with respect to $p = 1$. Indeed, as we increase p , the maximum difference consistently decreases until reaching $p = \infty$ (i.e., fully egalitarian), whose objective indeed is minimising such a maximum difference. Nonetheless, our results also show that the differences among the ethical principles for $p \geq 3$ are not very significant.

6.2 Multiple ethical principle case

In this second experiment we investigate how considering multiple ethical principles impacts the distribution of residuals, along the lines of the methodology followed in the previous section. Furthermore, in this case we are interested in evaluating whether our re-weighting approach achieves our goal of balancing the contributions of the different ethical principles in terms of the measure Ψ in Definition 3. To this end, we consider all the possible subsets of $\{1, 2, \infty\}$ as our set \mathbf{P} of ethical principles.⁴ Notice that our approach would allow us to consider any arbitrary set \mathbf{P} of ethical principles. Nonetheless, here we devote our attention to $p \in \{1, 2, \infty\}$ as these ethical principles can be semantically characterised in a clear way, as showed by the results of our previous experiment. Specifically, $p = 1$ results in a consensus that minimises the *average* residual, $p = 2$ to one that minimises the *standard deviation* of the residuals, and $p = \infty$ to one that minimises the *maximum* residual.

In particular, we are interested in evaluating the effect of the reweighting scheme proposed in Section 5 on the metrics that characterise the distribution of the residuals. Hence, in Table 1 we compare the descriptive statistics (namely the *average* and the *standard deviation*, as well as the *minimum* and *maximum*) of the distribution of the residuals obtained by solving Equation (21) (i.e., without reweighting) and Equation (23) (i.e., with reweighting).

We immediately observe that the consensus computed by solving Equation (21) for all the cases that include $p = 1$ (i.e., $\mathbf{P} = \{1, 2\}$, $\mathbf{P} = \{1, \infty\}$, and $\mathbf{P} = \{1, 2, \infty\}$) is exactly equal to the one computed for $p = 1$ in the single-norm case, experimentally confirming our discussion in Section 5. This is also quantitatively confirmed by the extremely high values of the measure Ψ , which indicates that the contributions of the different ethical principles are very unbalanced. Based on these results, we conclude that Equation (21) is not suitable to compute a consensus based on a set \mathbf{P} of multiple ethical principles including $p = 1$, since such a p dominates any other p .

On the other hand, we observe that the multi-norm consensus obtained with our reweighting scheme (Equation (23)) is characterised by descriptive statistics that are a balanced mix of its single components. In fact, as expected, the Ψ measure (i.e., the variance) is very close to zero, indicating that the contributions of the different ethical principles are practically equal. For instance, for $\mathbf{P} = \{1, 2\}$ we observe a lower σ^2 compared to $p = 1$ and a lower μ compared to $p = 2$. That is, we correctly observe the “contributions” of both $p = 1$ (i.e., reducing μ) and $p = 2$ (i.e., reducing σ^2). Similarly, for $\mathbf{P} = \{1, \infty\}$ we observe a lower maximum residual compared to $p = 1$ and a lower μ compared to $p = \infty$, and for $\mathbf{P} = \{2, \infty\}$ we observe a lower maximum residual compared to $p = 2$ and a lower σ^2 compared to $p = \infty$. Finally, when we combine all three ethical principles, we correctly obtain a consensus where μ is lower compared to $p = 2$ and $p = \infty$, σ^2 is lower compared to $p = 1$ and $p = \infty$, and maximum residual is lower compared to $p = 1$ and $p = 2$.

⁴We also consider subsets of size 1 (i.e., cases where \mathbf{P} contains only a single ethical principle), which will serve as a reference to better evaluate the multi-norm cases.

\mathbf{P}	Equation (21) (no reweighting)					Equation (23) (reweighting)				
	min	max	μ	σ^2	Ψ	min	max	μ	σ^2	Ψ
{1}	0.00	6.00	1.05	2.47	0.00	0.00	6.00	1.05	2.47	0.00
{2}	0.00	4.57	1.20	1.45	0.00	0.00	4.57	1.20	1.45	0.00
{ ∞ }	0.00	4.00	3.03	1.75	0.00	0.00	4.00	3.03	1.75	0.00
{1, 2}	0.00	6.00	1.05	2.47	6336.10	0.00	5.78	1.10	1.88	0.00
{1, ∞ }	0.00	6.00	1.05	2.47	7936.10	0.00	4.00	1.13	1.92	0.00
{2, ∞ }	0.00	4.00	1.21	1.45	85.93	0.00	4.00	1.21	1.45	0.00
{1, 2, ∞ }	0.00	6.00	1.05	2.47	6383.20	0.00	4.00	1.15	1.70	0.00

Table 1: Descriptive statistics for the residulas with all $\mathbf{P} \subseteq \{1, 2, \infty\}$. Notice that for single-norm cases (first 3 rows) Equation (21) and Equation (23) correctly provide the same results, since reweighting has no impact when \mathbf{P} contains only one p .

7 Discussion & Concluding Remarks

In this paper, we formulated and solved the problem of computing a consensus based on multiple ethical principles among the preferences of different individuals so that the ethical principles involved correctly contribute to such a consensus. Furthermore, we empirically analysed the impact of different configurations of single and multiple ethical principles on the resulting consensus.

The implications derived from this paper are twofold. From a theoretical perspective, using a p -metric as a proxy to represent a criterion or an ethical principle to reach a consensus allows us to extend previous analysis to account for multiple ethical principles that represent the actual diversity in individuals and social groups. From a practical perspective, we reformulate the problem to use a state-of-the-art algorithm that allows us to explore the characteristics of the possible consensus when multiple perspectives, not only the extreme values, are considered.

We should highlight that the choice of parameter p may be open to debate. Although p can be mathematically interpreted as the positioning between the principle of maximum freedom and the principle of maximum fairness, values of p different from 1, 2, and ∞ have a complex semantic interpretation. This conclusion has been possible thanks to our novel technique to compute the consensus for any value of p , including the intermediate ones (which, based on our knowledge of the state of the art, was not possible before). Along these lines, even if we could not provide a precise semantic interpretation of each ethical principle p , our contributions allowed us to conclude that any $p \geq 3$ is not significantly different from $p = \infty$. We also remark that our conclusions are bound to the particular dataset that we considered in our experiments, which might vary in other test cases. On the contrary, parameters λ are easier to interpret and assign. Since each value of λ stands for the social influence, or importance, of each particular individual (or group of people) involved in the group decision-making, it is easier to conduct sensitivity analysis or explore solutions compared to p .

Our work represents further motivation to enhance the integration of multiple ethical principles in group decision-making. First, it might be worth considering different distance functions other than p -metric distances (adopted both in this paper and in González-Pachón and Romero (2016)), as well as studying

the computational properties of the associated consensus computation problem. Second, our contribution allows us to compare our general approach to aggregate judgements with existing algorithms in the ranking aggregation literature to aggregate different rankings. This might spur interesting cross-fertilisation between these two different but related problems.

Finally, our work opens the door to *collective ethics*. As argued in Gabriel (2020), we live in a pluralistic world where people ascribe to different moral systems and therefore abide by different ethical principles. Making decisions that align with a group of people with different ethical principles poses a so-called *pluralistic value alignment* problem Gabriel (2020). Indeed, this is the case, for instance, when conducting policy-making decisions that align stakeholders with various ethical principles (e.g., Pigmans et al. (2017, 2019)). Indeed, as noted in Scharfbillig et al. (2021), policymakers must consider citizens' values (ethical principles) and identities when developing and communicating policies. We argue that the contributions in this paper make headway in this direction. The tools provided in this paper not only allow a group to reach a consensus but to do so while taking into account the collective ethics of the group. Future work should address the practical use of the tools presented here.

Acknowledgements

Research supported by projects: CI-SUSTAIN (PID2019-104156GB-I00); TAILOR (H2020-952215); 2021 SGR 00754 funded by Generalitat de Catalunya; VAE TED2021-131295B-C31, funded by MCIN/AEI/10.13 039/501100011033 and NextGenerationEU/ PRTR; and VALAWAI (Horizon Europe #101070930). Funding for open access charge: CRUE-Universitat Politècnica de València.

References

- R. S. Tindale, J. R. Winget, Group decision-making, in: Oxford Research Encyclopedia of Psychology, 2019.
- M. Yazdani, A. E. Torkayesh, P. Chatterjee, A. Fallahpour, M. J. Montero-Simo, R. A. Araque-Padilla, K. Y. Wong, A fuzzy group decision-making model to measure resiliency in a food supply chain: A case study in Spain, *Socio-Economic Planning Sciences* 82 (2022) 101257.
- M. Fagence, *Citizen participation in planning*, volume 19, Elsevier, 2014.
- C. L. Butler, A. Rothstein, *On conflict and consensus: A handbook on formal consensus decision making*, Takoma Park, 2006.
- Z. Xu, X. Cai, Group consensus algorithms based on preference relations, *Information Sciences* 181 (2011) 150–162.
- E. Herrera-Viedma, F. J. Cabrerizo, J. Kacprzyk, W. Pedrycz, A review of soft consensus models in a fuzzy environment, *Information Fusion* 17 (2014) 4–13.
- F. A. Ortega Riejos, M. Á. Pozo Montaña, J. Puerto Albandoz, Modelling and planning public cultural schedules for efficient use of resources, *Computers and Operations Research* 58 (2015) 9–23.

- I. J. Pérez, F. J. Cabrerizo, S. Alonso, Y. Dong, F. Chiclana, E. Herrera-Viedma, On dynamic consensus processes in group decision making problems, *Information Sciences* 459 (2018) 20–35.
- Z. Zhang, S.-M. Chen, C. Wang, Group decision making based on multiplicative consistency and consensus of fuzzy linguistic preference relations, *Information Sciences* 509 (2020) 71–86.
- Y. Lu, Y. Xu, E. Herrera-Viedma, Y. Han, Consensus of large-scale group decision making in social network: the minimum cost model based on robust optimization, *Information Sciences* 547 (2021) 910–930.
- A. Noori, H. Bonakdari, A. H. Salimi, B. Gharabaghi, A group multicriteria decision-making method for water supply choice optimization, *Socio-Economic Planning Sciences* 77 (2021) 101006.
- D. Pamucar, D. Macura, M. Tavana, D. Božanić, N. Knežević, An integrated rough group multicriteria decision-making model for the ex-ante prioritization of infrastructure projects: the serbian railways case, *Socio-Economic Planning Sciences* 79 (2022) 101098.
- Z. Zhang, S.-M. Chen, Group decision making based on multiplicative consistency and consensus of pythagorean fuzzy preference relations, *Information Sciences* 601 (2022) 340–356.
- P. Liu, K. Zhang, P. Wang, F. Wang, A clustering- and maximum consensus-based model for social network large-scale group decision making with linguistic distribution, *Information Sciences* 602 (2022) 269–297.
- X. Liang, J. Guo, P. Liu, A large-scale group decision-making model with no consensus threshold based on social network analysis, *Information Sciences* 612 (2022) 361–383.
- H. Hassani, R. Razavi-Far, M. Saif, E. Herrera-Viedma, Consensus-based decision support model and fusion architecture for dynamic decision making, *Information Sciences* 597 (2022) 86–104.
- H. Zhang, F. Wang, Y. Dong, F. Chiclana, E. Herrera-Viedma, Social trust driven consensus reaching model with a minimum adjustment feedback mechanism considering assessments-modifications willingness, *IEEE Transactions on Fuzzy Systems* 30 (2021) 2019–2031.
- J. Xiao, X. Wang, H. Zhang, Exploring the ordinal classifications of failure modes in the reliability management: An optimization-based consensus model with bounded confidences, *Group Decision and Negotiation* 31 (2022) 49–80.
- J. González-Pachón, C. Romero, Distance-based consensus methods: a goal programming approach, *Omega* 27 (1999) 341–347.
- J. González-Pachón, C. Romero, Aggregation of ordinal and cardinal preferences: a framework based on distance functions, *Journal of Multi-Criteria Decision Analysis* 15 (2008) 79–85.
- J. González-Pachón, C. Romero, The design of socially optimal decisions in a consensus scenario, *Omega* 39 (2011) 179–185.

- J. González-Pachón, C. Romero, Bentham, Marx and Rawls ethical principles: In search for a compromise, *Omega* 62 (2016) 47–51.
- J. Bentham, *An introduction to the principles of morals and legislation*, Payne and Son, London, 1789.
- J. Rawls, *A theory of justice*, Oxford University Press, Oxford, 1973.
- S. Boyd, L. Vandenberghe, *Convex optimization*, Cambridge university press, 2004.
- D. Adil, R. Peng, S. Sachdeva, Fast, provably convergent IRLS algorithm for p-norm linear regression, in: *Advances in Neural Information Processing Systems*, 2019, pp. 14189–14200.
- C. Dwork, R. Kumar, M. Naor, D. Sivakumar, Rank aggregation methods for the web, in: *Proceedings of the 10th international conference on World Wide Web*, 2001, pp. 613–622.
- J. Ding, D. Han, J. Dezert, Y. Yang, A new hierarchical ranking aggregation method, *Information Sciences* 453 (2018) 168–185.
- J. A. Aledo, J. A. Gámez, A. Rosete, A highly scalable algorithm for weak rankings aggregation, *Information Sciences* 570 (2021) 144–171.
- W. D. Cook, T. Raviv, A. J. Richardson, Aggregating incomplete lists of journal rankings: An application to academic accounting journals, *Accounting perspectives* 9 (2010) 217–235.
- B. Fishbain, E. Moreno-Centeno, Self calibrated wireless distributed environmental sensory networks, *Scientific reports* 6 (2016) 1–10.
- E. Moreno-Centeno, A. R. Escobedo, Axiomatic aggregation of incomplete rankings, *IIE Transactions* 48 (2016) 475–488.
- J. G. Kemeny, L. Snell, Preference ranking: an axiomatic approach, *Mathematical models in the social sciences* (1962) 9–23.
- Y. Yoo, A. R. Escobedo, A new binary programming formulation and social choice property for kemeny rank aggregation, *Decision Analysis* 18 (2021) 296–320.
- A. R. Escobedo, E. Moreno-Centeno, R. Yasmin, An axiomatic distance methodology for aggregating multimodal evaluations, *Information Sciences* (2022).
- C. Romero, Extended lexicographic goal programming: a unifying approach, *Omega* 29 (2001) 63–71.
- D. Jones, M. Tamiz, et al., *Practical goal programming*, volume 141, Springer, New York, 2010.
- A. B. Atkinson, et al., On the measurement of inequality, *Journal of Economic Theory* 2 (1970) 244–263.

- A. Charnes, W. W. Cooper, Goal programming and multiple objective optimizations: Part 1, *European Journal of Operational Research* 1 (1977) 39–54.
- R. Caballero, F. Ruiz, M. V. R. Uría, C. Romero, Interactive meta-goal programming, *European Journal of Operational Research* 175 (2006) 135–154.
- T. L. Saaty, A scaling method for priorities in hierarchical structures, *Journal of Mathematical Psychology* 15 (1977) 234–281.
- M. Samejima, Y. Matsushita, Fast general norm approximation via iteratively reweighted least squares, in: *Asian Conference on Computer Vision*, Springer, 2016, pp. 207–221.
- F. Alizadeh, D. Goldfarb, Second-order cone programming, *Mathematical Programming* 95 (2001).
- E.-M. Nordström, C. Romero, L. O. Eriksson, K. Öhman, Aggregation of preferences in participatory forest planning with multiple criteria: an application to the urban forest in lycksele, sweden, *Canadian Journal of Forest Research* 39 (2009) 1979–1992.
- I. Gabriel, Artificial intelligence, values, and alignment, *Minds and Machines* 30 (2020) 411–437.
- K. Pigmans, N. Doorn, H. Aldewereld, V. Dignum, Decision-making in water governance: From conflicting interests to shared values, in: *Responsible Innovation*, Springer, 2017, pp. 165–178.
- K. Pigmans, H. Aldewereld, V. Dignum, N. Doorn, The role of value deliberation to improve stakeholder participation in issues of water governance, *Water Resources Management* 33 (2019) 4067–4085.
- M. Scharfbillig, L. Smillie, D. Mair, M. Sienkiewicz, J. Keimer, R. Pinho Dos Santos, H. Vinagreiro Alves, E. Vecchione, S. L., Values and Identities - a policymaker’s guide – Executive summary, Technical Report, Publications Office of the European Union, 2021.